

Affine Dimers from Characteristic Polygons

D. HOLMES*

Abstract - Recent work by Forsgård indicates that not every convex lattice polygon arises as the characteristic polygon of an affine dimer or, equivalently, an admissible oriented line arrangement on the torus in general position. We begin the classification of convex lattice polygons arising as characteristic polygons of affine dimers. We present several general constructions of new affine dimers from old, and an algorithm for finding affine dimers with prescribed polygon.

With these tools we prove that all lattice triangles, generalised parallelograms, and polygons of genus at most two admit an affine dimer.

Keywords : dimer model; hyperplane arrangement; torus; lattice polygon

Mathematics Subject Classification (2020) : 52C35; 52B20; 05B45

1 Introduction

A dimer model is an embedded bipartite graph on the torus \mathbb{T}^2 or, depending on the application, any surface Σ . They were originally introduced in statistical mechanics to model molecular interactions. In a simplified model, the thermodynamic properties of a mixture of molecules can be calculated from a combinatorial factor that counts the number of arrangements of molecules on a square lattice. If all molecules are dimers rather than monomers or higher polymers, this amounts to counting the number of domino tilings of the square lattice [12]. Thinking of the square lattice as an embedded graph, there is a one-to-one correspondence between domino tilings of the lattice and perfect matchings of the graph. The requirement that the graph is bipartite arises when one takes into account the two possible charges of a particle. Finally, the torus \mathbb{T}^2 is a natural choice of ambient space to account for translational symmetries such as that of a crystal. More recent applications of dimer models can be found in algebraic and tropical geometry, as well as in string theory (e.g., [16] and [14]).

To every dimer model on \mathbb{T}^2 one can associate a convex lattice polygon, called the *characteristic polygon*, in at least two ways. The first way is as the Newton polygon of the determinant of the Kasteleyn operator, a generalisation of the adjacency matrix where the entries are weighted according to their meridional and longitudinal winding numbers ([3], Section 7). A variant of this operator was used by Kasteleyn in [12] to calculate the

*This work was supported by the London Mathematical Society and the Department of Pure Mathematics and Mathematical Statistics, University of Cambridge.



number of domino tilings of a rectangular square lattice as a Pfaffian. The second way is as the convex hull of the values of the height function, which assigns a value in \mathbb{Z}^2 to every perfect matching of the dimer model. These two notions turn out to be equivalent, and it is natural to consider the inverse problem: For which convex lattice polygons does there exist a dimer model with characteristic polygon as prescribed?

This question has been answered positively for all convex polygons if no further restrictions are imposed on the dimer model [8]. Futaki–Ueda and Ueda–Yamzaki realised that, in some cases, the dimer model may be obtained from the faces of a certain hyperplane arrangement on \mathbb{T}^2 ([7], [19], [20], as cited in [5]). In this case, the dimer is called *affine*. However, Forsgård exhibited a family of convex polygons which do not admit an affine dimer ([5], Section 4).

The goal of this paper is to classify which convex polygons admit an affine dimer. We present partial results consisting of a list of constructions to obtain new dimers from old, an algorithm implemented in Java to verify whether a polygon admits an affine dimer, and a positive answer for all convex lattice polygons that are triangles, “generalised parallelograms”, or have at most two interior lattice points. These results are summarised in Theorem A & B at the end of this section.

The results have the following application to algebraic and tropical geometry. Given a complex curve C in $(\mathbb{C}^*)^2$, the *coamoeba* $\mathcal{C} \subseteq \mathbb{T}^2$ is its image under the argument projection $(x, y) \mapsto (\arg(x), \arg(y))$ which naturally takes values on \mathbb{T}^2 . The *shell* of the coamoeba is a line arrangement \mathcal{H} on \mathbb{T}^2 that is derived from the bivariate polynomial defining C and satisfies $\mathcal{C} = \mathcal{C} \cup \mathcal{H}$ (c.f. [11] and [5]). Then \mathcal{H} divides \mathbb{T}^2 into several tiles, and we say that a tile is *full* if it is fully contained in \mathcal{C} . We say that the coamoeba is represented by a dimer if we can embed a bipartite graph on \mathbb{T}^2 such that every vertex is contained in the interior of a full tile, every tile contains at most one vertex, and the edges correspond to shared corners between two tiles. If such a graph exists, it is by definition a dimer, and automatically affine since it comes from the line arrangement \mathcal{H} .

It is natural to ask which complex curves possess coamoebas that are represented by a dimer. An important observation is that the Newton polygon of the defining polynomial of the curve is the same as the characteristic polygon of the affine dimer representing its coamoeba, if such a dimer exists. Therefore, a first obstruction is the non-existence of an affine dimer with given characteristic polygon. We prove that this combinatorial obstruction vanishes if the genus of the curve is at most two.

Our results also imply that all tropical curves of genus ≤ 2 can be lifted to an exact Lagrangian submanifold of $(\mathbb{C}^*)^2$, as described in [9].

For simplicity, we work with *homology polygons* rather than *characteristic polygons* from height functions. These concepts are equivalent, as outlined in the very readable source [2].

1.1 Definitions

The n -dimensional torus \mathbb{T}^n is the quotient $\mathbb{R}^n/\mathbb{Z}^n$ with quotient map $q : \mathbb{R}^n/\mathbb{Z}^n \rightarrow \mathbb{T}^n$. Note that q is a universal cover for \mathbb{T}^n .



Definition 1.1 A dimer is a bipartite multigraph $G = (V_\circ \sqcup V_\bullet, E)$ embedded on the two-dimensional torus. (This means we allow multiple edges between two vertices).

Let $\hat{H} \subseteq \mathbb{R}^n$ be an affine hyperplane, i.e., there exist $a \in \mathbb{R}^n, b \in \mathbb{R}$ such that $\hat{H} = \{x \in \mathbb{R}^n : \langle a, x \rangle = b\}$. We call $H := q(\hat{H})$ a *hyperplane on the torus*. If $\dim(\hat{H}) = 1$, we call H a *line (on the torus)*.

We now specialise to $n = 2$. A *closed geodesic* is a closed loop given by a line $H \subseteq \mathbb{T}^2$. Once we fix a choice of orientation for a closed geodesic H , there are unique coprime integers $a, b \in \mathbb{Z}$ such that the homology class of H is $[H] = (a, b) \in H_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, i.e., (a, b) is the direction of H . We call $(a, b) \in \mathbb{Z}^2$ *primitive* if $\gcd(a, b) = 1$. One can interpret a and b as the winding numbers of H around the two directions of $\mathbb{T}^2 \cong \mathbb{T} \times \mathbb{T}$. Note, however, that this choice of directions is by no means intrinsic to \mathbb{T}^2 and we will exploit this symmetry on several occasions, using the action by automorphisms of $GL_2(\mathbb{Z})$ on \mathbb{T}^2 .

Definition 1.2 An *oriented line arrangement (on the 2-torus)* is a finite set \mathcal{H} of closed geodesics on \mathbb{T}^2 . The line arrangement is called

- *in general position* if no three lines intersect in a point, parallel lines are disjoint, and not all lines are parallel;
- *admissible* if it is in general position and every oriented line segment is a boundary component of a face whose edges are consistently oriented, i.e., all clockwise or all counterclockwise. (A line segment of \mathcal{H} is a segment of a line in \mathcal{H} whose endpoints are intersection points of \mathcal{H} and whose interior contains no intersection points.)

Note that if a line arrangement is in general position then all faces are automatically homeomorphic to a disk. Figure 1 gives an example of admissible and non-admissible line arrangements. Note that these only differ by a translation of the upper horizontal line, so both arrangements represent the same multiset of homology classes in $H_1(\mathbb{T}^2)$.

We briefly elaborate on the equivalence of admissible oriented line arrangements and a certain class of dimers, called *affine dimers*. Given an admissible oriented line arrangement, we obtain a dimer $G = (V_\circ \sqcup V_\bullet, E)$ as follows. Let V_\circ and V_\bullet be the sets of faces oriented clockwise and counterclockwise, respectively. For each intersection point of the line arrangement, we add an edge to E connecting the two oriented faces meeting there. The obtained graph is bipartite. To embed G in \mathbb{T}^2 we place a vertex in the interior of each consistently oriented face. Each edge can then be realised as a union of two line segments meeting at the shared intersection point of the two faces (see Figure 2).



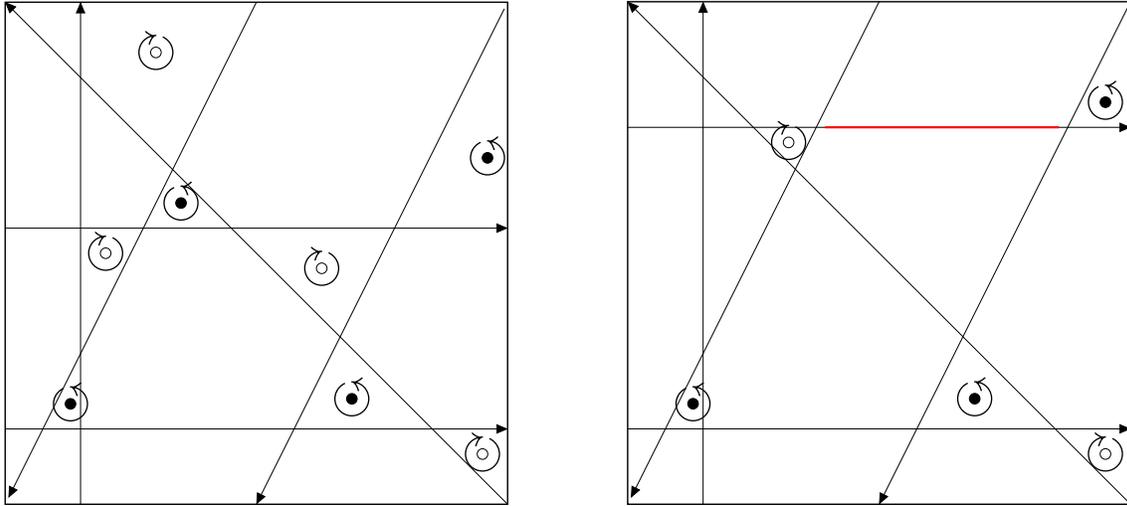


Figure 1: Examples of an admissible (left) and non-admissible (right) oriented line arrangement. Consistently oriented faces are indicated with \ominus and \oplus . The example to the right is not admissible because, for example, the red line segment does not bound a consistently oriented face.

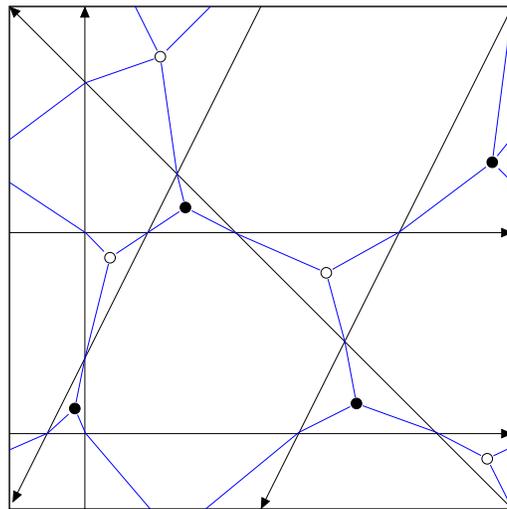


Figure 2: The affine dimer $G = (V_{\circ} \sqcup V_{\bullet}, E)$ obtained from the admissible line arrangement in Figure 1. The edges of G are depicted in blue.

The converse construction is also possible: Given a dimer $G = (V_{\circ} \sqcup V_{\bullet}, E)$ such that

- the vertices of G are faces of a line arrangement in general position;
- each edge of G connects two faces along an intersection point of the line arrangement such that the connected faces are opposite each other at that intersection point;



- each intersection point of the line arrangement lies on exactly one edge of G and each edge of G contains exactly one intersection point of the line arrangement,

then we may declare the faces in V_\circ and V_\bullet to be oriented clockwise and counterclockwise, respectively. The above conditions determine a well-defined choice of orientation for each line. Thus, we obtain an admissible oriented line arrangement.

Definition 1.3 *An affine dimer is a dimer satisfying the three conditions above. Figure 2 gives an example.*

Thus, the notions of an affine dimer and an admissible oriented line arrangement are equivalent, and we will use them interchangeably.

1.2 Problem Statement

Our leading question is the following:

Question 1.4 *For which multisets of homology classes $S = \{h_1, \dots, h_n\} \subset H_1(\mathbb{T}^2) \cong \mathbb{Z}^2$ is there an admissible oriented line arrangement $\mathcal{H} = \{H_1, \dots, H_n\}$ whose lines represent S , i.e., such that $[H_i] = h_i$ for $i = 1, \dots, n$?*

Figure 1 shows that a multiset of homology classes may be represented both by admissible and non-admissible oriented line arrangements. Moreover, Forsgård showed that there is a family of multisets of homology classes indexed by $\mathbb{N}_{\geq 5}$, for which there are no admissible oriented line arrangements representing them [5]. Thus, the problem is non-trivial.

We already saw that the homology class of a closed geodesic on \mathbb{T}^2 is automatically *primitive*, i.e., $(a, b) \in \mathbb{Z}^2$ with $\gcd(a, b) = 1$. There is another immediate necessary condition, which will allow us to reformulate the problem in terms of convex polygons on the integer lattice.

Lemma 1.5 *Let $\mathcal{H} = \{H_1, \dots, H_n\}$ be an admissible oriented line arrangement representing the homology classes $[H_i] = (a_i, b_i) \in \mathbb{Z}^2$. Then*

$$\sum_{i=1}^n [H_i] = 0.$$

Proof. Each oriented line H_i is subdivided into several oriented line segments whose endpoints are intersection points of the line arrangement. These oriented line segments represent 1-chains on \mathbb{T}^2 , so we may write $H_i = \sum_{\text{segments } e \text{ of } H_i} e$, where $e \in C_1(\mathbb{T}^2)$ is a line segment of H_i . On the other hand, each segment belongs to exactly one consistently oriented face. Thus, as chains

$$\sum_{i=1}^n H_i = \sum_{i=1}^n \left(\sum_{\text{segments } e \text{ of } H_i} e \right) = \sum_{\text{segments } e \text{ of } \mathcal{H}} e = \sum_{\text{oriented faces } F} \left(\sum_{\text{edges } e \text{ of } F} e \right).$$



Passing to homology, we get $\left[\sum_{\text{edges } e \text{ of } F} e\right] = 0$ for each oriented face F . Therefore $\sum_{i=1}^n [H_i] = 0$. \square

Definition 1.6 A lattice polygon is a polygon in \mathbb{R}^2 whose vertices all lie in \mathbb{Z}^2 .

Lemma 1.7 There is a bijection between the finite multisets of primitive elements of \mathbb{Z}^2 summing to zero and the convex lattice polygons on \mathbb{Z}^2 up to translation.

Proof. Given primitive elements $h_i = (a_i, b_i) \in \mathbb{Z}^2$, i.e., $\gcd(a_i, b_i) = 1$, we may order them by their angle $\arg(h_i) \in [0, 2\pi)$ with the x -axis. We define a convex lattice polygon via the vertices $v_0 = (0, 0)$ and $v_i = v_{i-1} + h_i$. This is a closed polygon since $\sum_{i=1}^n h_i = 0$ and convex since we ordered the h_i .

Conversely, given a convex lattice polygon, orient the edges counterclockwise and subdivide each edge so that it contains no integer lattice point in its interior. Viewing each edge as a vector $(a, b) \in \mathbb{Z}^2$, this corresponds exactly to $\gcd(a, b) = 1$, and thus we obtain a primitive element $h_i = (a_i, b_i) \in \mathbb{Z}^2$ for each primitive edge segment of the polygon. Finally, $\sum_{i=1}^n h_i = 0$ since polygons are closed. \square

Definition 1.8 Let \mathcal{H} be an admissible oriented line arrangement on \mathbb{T}^2 . The homology polygon P of \mathcal{H} is the convex lattice polygon obtained from the homology classes of the lines of \mathcal{H} . Equivalently, we can talk about the homology polygon of an affine dimer. Figure 3 shows an example.

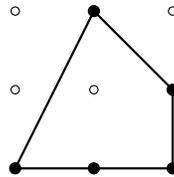


Figure 3: Homology polygon of the admissible line arrangement in Figure 1 and the affine dimer in Figure 2. The homology classes represented by the polygon are $(1, 0), (1, 0), (0, 1), (-1, 1), (-1, -2)$.

It is not a priori clear that the homology polygon of an affine dimer is well-defined, as there might exist different admissible line arrangements with different homology classes that give the same affine dimer $G = (V_\circ \sqcup V_\bullet, E)$ via the construction in Section 1.1. However, as mentioned earlier, the homology polygon is equivalent to the *characteristic polygon* (see [2]) which only depends on the data of the dimer and not the line arrangement. Hence, the homology polygon of an affine dimer is well-defined. We only prefer to use homology polygons in the problem statement because they are slightly easier to define.

Thus, we can reformulate our question:

Question 1.9 (Reformulation of Question 1.4) Which convex lattice polygons arise as the homology polygon of an admissible oriented line arrangement? Or which convex lattice polygons admit an affine dimer?



1.2.1 Invariance under $GL_2(\mathbb{Z})$

Finally, we note that $GL_2(\mathbb{Z}) = \{A \in \mathbb{Z}^{2 \times 2} : \det(A) = \pm 1\}$ acts on \mathbb{T}^2 by linear automorphisms which preserve admissible oriented line arrangements. Similarly, $GL_2(\mathbb{Z})$ acts on the space of convex lattice polygons through its action on \mathbb{Z}^2 , preserving area and the number of lattice points in the interior and on the boundary. We call two lattice polygons *equivalent* if they are related by an action of $GL_2(\mathbb{Z})$ and translation by a vector in \mathbb{Z}^2 . Thus, whether a convex lattice polygon admits an affine dimer only depends on its equivalence class, and we arrive at our final formulation of the problem:

Question 1.10 (*Reformulation of Question 1.9*) *Which equivalence classes of convex lattice polygons arise as the homology polygon of an admissible oriented line arrangement? Or which equivalence classes admit an affine dimer?*

1.3 Outline of Results and Structure

Section 2 surveys some basic combinatorial properties of affine dimers and motivates the name *genus* for the number of interior points of the homology polygon P , by connecting it to the genus of a punctured compact orientable surface that is homotopy equivalent to G .

In Section 3 we present three constructions of affine dimers. The “double everything”-construction exhibits an affine dimer for every lattice polygon consisting of pairs of antiparallel primitive side segments. The other two constructions (“lifting” and “adding an antiparallel pair”) give new dimers from old:

Theorem A *Let P be the homology polygon of an affine dimer.*

- (i) *If $B \in \mathbb{Z}^{2 \times 2}$ and $\det(B) \neq 0$ then $B(P)$ is also the homology polygon of an affine dimer.*
- (ii) *If $h \in \mathbb{Z}^2$ is a primitive side segment of P , then P_h is also the homology polygon of an affine dimer, where P_h is obtained from P by adding the side segments $h, -h$.*

Proof. (i) Proposition 3.1. (ii) Corollary 3.4 below. □

Section 4 summarises our algorithms, including a description of the moduli space $\mathcal{M} \cong \mathbb{T}^n$ of line arrangements representing a given homology polygon P .

Finally, Section 5 connects these results to finish the proof of Theorem B:

Theorem B *Let P be a convex lattice polygon such that*

- (i) *P is a triangle, or*
- (ii) *the primitive side segments of P are pairs of antiparallel side segments, or*
- (iii) *the number of interior lattice points of P is at most 2.*

Then P admits an affine dimer.

Proof. (i) Proposition 5.1. (ii) Proposition 3.2. (iii) Propositions 5.2, 5.4, 5.5 below. □



2 Basic Combinatorics of Affine Dimers

In this section we develop some basic combinatorics of affine dimers.

Definition 2.1 Let $G = (V_\circ \sqcup V_\bullet, E)$ be an affine dimer with corresponding admissible oriented line arrangement $\mathcal{H} = \{H_1, \dots, H_n\}$. Then we denote by

- n ... the number of lines,
- $f_\circ := |V_\circ|, f_\bullet := |V_\bullet|$... the number of faces of the line arrangement oriented clockwise and anticlockwise, respectively,
- f_\times ... the number of faces that are inconsistently oriented,
- $f = f_\circ + f_\bullet + f_\times$... the number of faces of the line arrangement,
- v ... the total number of vertices of the line arrangement, i.e., intersection points of lines in \mathcal{H} ,
- e_\circ, e_\bullet ... the number of line segments of \mathcal{H} belonging to faces in V_\circ or V_\bullet , respectively,
- $e = e_\circ + e_\bullet$... the number of line segments of the line arrangement \mathcal{H} ,
- g ... the genus of the dimer, which will be introduced in Section 2.2.

For example, the affine dimer in Figure 2 has $n = 5, f_\circ = f_\bullet = 4, f_\times = 5, f = v = e/2 = 13, e_\circ = e_\bullet = 13,$ and $g = 1$.

Proposition 2.2 (Basic counting)

$$(i) \ v - e + f = 0 \quad (ii) \ e_\circ = e_\bullet \quad (iii) \ f_\circ = f_\bullet \quad (iv) \ v = f = e/2$$

Proof. The proofs are as follows:

- (i) Immediate since \mathbb{T}^2 has Euler characteristic zero and a line arrangement in general position gives a CW decomposition of \mathbb{T}^2 .
- (ii) Each of the n closed geodesics consists alternately of edges counted by e_\circ and e_\bullet .
- (iii) This follows from the existence of a perfect matching for G (see Proposition 2.3).
- (iv) By (i) it suffices to show $v = f$, for which we induct on the number of lines. Adding a closed geodesic in general position adds as many faces as it adds vertices. To verify the induction basis, assume we only have two closed geodesics which are not parallel. As the homology class of a closed geodesic is a primitive element of \mathbb{Z}^2 , by the Euclidean algorithm we may use an action of $SL_2(\mathbb{Z})$ to assume that the geodesics have homology classes $(1, 0)$ and (c, d) . By inspection, this configuration has $v = f = d$.



□

Proposition 2.3 *An affine dimer G admits a perfect matching.*

Proof. See [2] for a very readable discussion of the perfect matchings of a (not necessarily affine) dimer, whose information is encoded in the *characteristic polygon* via height functions. This is a special case.

Let $\rho \in \mathbb{R}^2 \setminus \{0\}$ be a vector that does not indicate the (signed) direction of any line in \mathcal{H} . Then every consistently oriented face has a vertex at which the directions of the two intersecting lines are immediately to the right and to the left of ρ . Now match each clockwise face in V_\circ to the counterclockwise face in V_\bullet adjacent to it via that vertex. This defines a bijection $V_\circ \rightarrow V_\bullet$ whose inverse $V_\bullet \rightarrow V_\circ$ is constructed identically using the same ρ . Thus, we have a matching.

This construction is illustrated in Figure 4. □

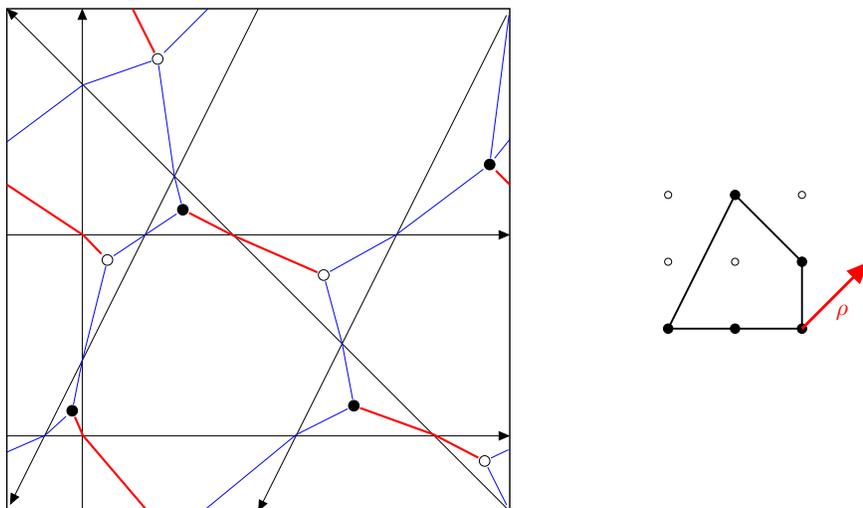


Figure 4: Matching (red) of the affine dimer in Figure 2 corresponding to $\rho = (1, 1)$ (left). As seen in the dimer's homology polygon (right), any ρ with $\arg(\rho) \in (0, \pi/2)$ produces the same matching.

Proposition 2.4 *We have*

$$1/3 \leq f_\times/f \leq 1/2$$

with $f_\times/f = 1/2$ if and only if every inconsistently oriented face is a 4-gon, and $f_\times/f = 1/3$ if and only if every consistently oriented face is a triangle.

Proof. For the upper bound we count the number of corners of inconsistently oriented faces in two ways. On the one hand, this is $2v$ since each vertex is incident with two inconsistently oriented faces. On the other hand, each inconsistently oriented face has an



even number of vertices, as otherwise G contains an odd cycle contradicting bipartiteness. Thus,

$$2v \geq 4f_{\times}.$$

By Proposition 2.2, $v = f$, so the upper bound follows with equality condition as desired.

For the lower bound we count the number of corners of consistently oriented faces in two ways. Again, this is $2v$. But every face has at least three edges, so

$$2v \geq 3(f_{\circ} + f_{\bullet}) = 3(f - f_{\times}).$$

The result follows using $f = v$ again. □

2.1 Area of the Homology Polygon

The next result follows directly from Johansson's and Forsgård's work in [6].

Theorem 2.5 *Let P be the homology polygon of an affine dimer. Then*

$$f_{\times} = 2\text{Area}(P).$$

Proof sketch. The key step is to show that $2\pi\text{Area}(P)$ is the sum of *inner angles* of vertices of the line arrangement, counting one per vertex ([6], Lemma 3.2). Here, the inner angle of two oriented intersecting lines is defined to be positive and lies between an ingoing and an outgoing ray.

Thus, $2\pi\text{Area}(P)$ is the sum of interior angles of all clockwise faces. This is exactly half the sum of exterior angles of all inconsistently oriented faces, which is 2π per face. Thus,

$$2\pi\text{Area}(P) = \frac{2\pi f_{\times}}{2}.$$

□

For example, the affine dimer in Figure 4 has $f_{\times} = 5 = 2\text{Area}(P)$.

2.2 Genus of an Affine Dimer

We now describe two ways to think of an affine dimer (or equivalently of an admissible oriented line arrangement) as a two-dimensional geometric shape.

Definition 2.6 *The realisation of an affine dimer G is the set $\bar{G} := \bigcup_{F \in V_{\circ} \sqcup V_{\bullet}} \bar{F} \subseteq \mathbb{T}^2$, i.e., the union of the closed oriented faces of the admissible oriented line arrangement \mathcal{H} .*

By definition, this depends on the choice of admissible line arrangement \mathcal{H} corresponding to G . However, many properties of \bar{G} only depend on the homology polygon P .

Proposition 2.7 *The Euler characteristic of \bar{G} is $\chi(\bar{G}) = -f_{\times}$.*



Proof.

$$\chi(\bar{G}) = v - e + (f_{\circ} + f_{\bullet}) = v - e + f - f_{\times} = \chi(\mathbb{T}^2) - f_{\times} = -f_{\times}.$$

□

The embedding of G described in Section 1.1 is a deformation retract of \bar{G} , so $\chi(G) = \chi(\bar{G})$.

Corollary 2.8 $\chi(G) = \chi(\bar{G}) = -2\text{Area}(P)$.

We may consider \bar{G} as the projection of a punctured smooth compact oriented surface \hat{G} embedded in \mathbb{R}^3 . To this end we use the smooth standard embedding $\varphi : \mathbb{T}^2 \hookrightarrow \mathbb{R}^3$ and consider $\varphi(\bar{G}) \subset \mathbb{R}^3$.

Definition 2.9 (The smooth orientable surface $\hat{G} \subset \mathbb{R}^3$) *Away from intersection points of boundary components of $\varphi(\bar{G})$ we identify \hat{G} with $\varphi(\bar{G})$. Near an intersection point, we exploit the third dimension and let \hat{G} twist locally by 180° like a helicoid as shown in Figure 5, i.e., the normal vector changes smoothly from v to $-v$ when traversing this neighbourhood along an edge of G . These patches are glued together using bump functions so that we obtain a smooth compact embedded surface $\hat{G} \subset \mathbb{R}^3$.*

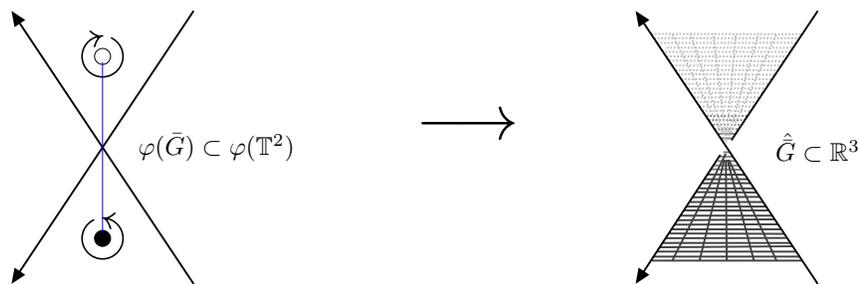


Figure 5: Twisting of \hat{G} in \mathbb{R}^3 near an intersection point of the boundary components of $\varphi(\bar{G})$. \hat{G} looks locally like the 180° segment of a helicoid near such points.

Proposition 2.10 \hat{G} is orientable.

Proof. An orientation $N : \hat{G} \rightarrow S^2$ is obtained as follows. By the Jordan–Brouwer separation theorem, $\mathbb{R}^3 \setminus \varphi(\mathbb{T}^2)$ consists of a bounded and an unbounded component. Away from intersection points of boundary components of $\varphi(\bar{G})$, let $N(p)$ point into the unbounded component at p if $p \in \varphi(\bigcup V_{\bullet})$ and into the bounded component if $p \in \varphi(\bigcup V_{\circ})$. Near an intersection point, let N twist as prescribed by the local helicoid in Figure 5. These local definitions of N glue together to form a well-defined orientation of \hat{G} because $G = (V_{\circ} \sqcup V_{\bullet}, E)$ is bipartite, so G does not contain a circuit of odd length. □

Proposition 2.11 \bar{G} and \hat{G} are homotopy equivalent. Therefore,

$$\chi(\hat{G}) = \chi(\bar{G}) = \chi(G) = -f_{\times} = -2\text{Area}(P).$$



Proof. It suffices to show that $\varphi(\bar{G})$ and \hat{G} are homotopy equivalent. Away from intersection points of boundary components of $\varphi(\bar{G})$, both surfaces are identical. Near an intersection point, the surfaces are equivalent by the homotopy that projects the right hand side of Figure 5 onto the left hand side. These homotopies glue together compatibly and hence $\varphi(\bar{G}) \simeq \hat{G}$. The equalities now follow from Proposition 2.7 and Corollary 2.8. \square

Lemma 2.12 (Pick's formula) *Let P be a simple lattice polygon (i.e., ∂P does not self-intersect and has exactly one connected component). Then*

$$\text{Area}(P) = |\overset{\circ}{P} \cap \mathbb{Z}^2| + \frac{1}{2} |\partial P \cap \mathbb{Z}^2| - 1.$$

Proof. This is a well-known result with many different proofs available. E.g., one standard proof is via Euler's formula [1], while a more non-standard proof uses the Weierstraß \wp -function [4]. \square

Theorem 2.13 *\hat{G} is homeomorphic to the compact oriented surface $\Sigma_{g,n}$ obtained by removing n disjoint open discs from the compact oriented surface Σ_g of genus g without boundary. Moreover, the genus g is the number of interior points of P , i.e.,*

$$\hat{G} \cong \Sigma_{g,n} \quad \text{where} \quad g = |\overset{\circ}{P} \cap \mathbb{Z}^2|.$$

Proof. The first part follows from the classification of surfaces and the fact that \hat{G} has n boundary components, one for each line in \mathcal{H} . Adding a puncture to a surface (i.e., removing an open disc) decreases the Euler characteristic by one. Thus, by Proposition 2.11,

$$\chi(\Sigma_g) = \chi(\hat{G}) + n = -2\text{Area}(P) + n.$$

But n is the number of primitive side segments of P which equals the number of lattice points on the boundary ∂P . Using $\chi(\Sigma_g) = 2 - 2g$ we get

$$g = 1 - \chi(\Sigma_g)/2 = 1 + \text{Area}(P) - |\partial P \cap \mathbb{Z}^2|/2.$$

Now the statement follows immediately from Lemma 2.12 (Pick's formula). \square

This explains our definition of the genus of a dimer:

Definition 2.14 *The genus of an affine dimer G with homology polygon P is $g := |\overset{\circ}{P} \cap \mathbb{Z}^2|$, the number of lattice points in the interior of P . We also call this the genus of the convex lattice polygon P .*

This matches the relation between the genus of a tropical curve in \mathbb{R}^2 and the number of interior points of its Newton polygon [15].

For example, Figure 6 shows an affine dimer G of genus zero. Since its line arrangement has $n = 3$ lines, $\hat{G} \cong \Sigma_{0,3}$ is the 3-punctured sphere also known as *pair of pants*. The affine dimer in Figure 4 has genus one and $\hat{G} \cong \Sigma_{1,5}$, the 5-punctured torus.



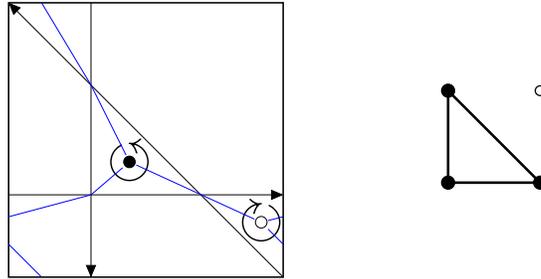


Figure 6: An affine dimer of genus zero and its homology polygon.

It is shown in Section 5 that every convex lattice polygon of genus at most 2 is the homology polygon of an affine dimer, answering Question 1.10 positively for these polygons. Moreover, by Proposition 5.1, every lattice triangle admits an affine dimer. Since for every $g \in \mathbb{N}$ there exists a lattice triangle of genus g , there exist affine dimers of all genera.

3 Constructions of Affine Dimers

Next, we present three constructions of affine dimers and analyse the obtained homology polygons.

3.1 Adding parallel edges

Proposition 3.1 *Let \mathcal{H} be an admissible oriented line arrangement with homology polygon P . Let $h \in \mathbb{Z}^2$ be a primitive side segment of P . Then the convex lattice polygon P_h obtained by adding the antiparallel side segments h and $-h$ to P is the homology polygon of an admissible oriented line arrangement. Thus, if P admits an affine dimer then so does P_h .*

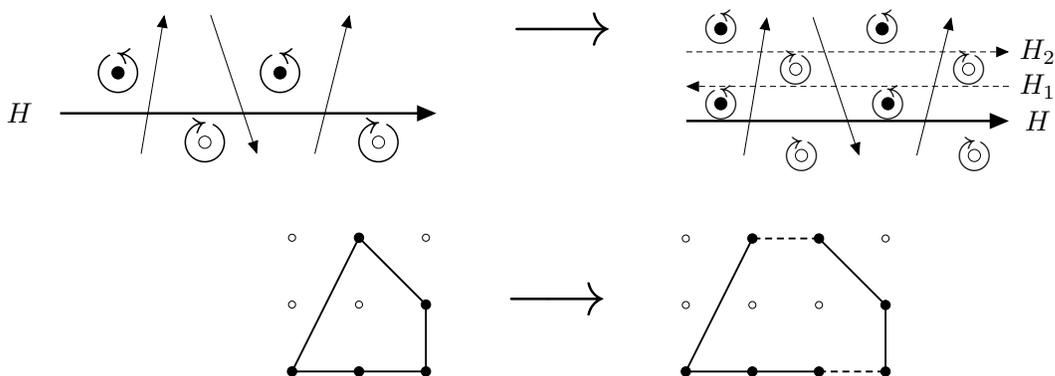


Figure 7: Constructing \mathcal{H}_H from \mathcal{H} (top) and P_h from P (bottom) for $h = (1, 0)$. Only the local picture near H is displayed, as everything else remains unchanged.



Proof. Let $H \in \mathcal{H}$ with $[H] = h$. We construct a new admissible oriented line arrangement \mathcal{H}_H with homology polygon P_h by adding two antiparallel lines H_1, H_2 with $[H_1] = -h$ and $[H_2] = h$. As depicted in Figure 7, we place them in the order H_2, H_1, H and close enough to H so that no other intersection points of \mathcal{H} lie between H_2 and H .

If k is the number of intersection points on H then this construction adds k consistently oriented faces to \mathcal{H} locally near H , $k/2$ of each orientation. Away from H the arrangement remains unchanged. Thus, we have obtained a new admissible arrangement \mathcal{H}_H of homology polygon P_h , as required. \square

3.2 Double everything

All lattice polygons consisting of pairwise antiparallel primitive side segments admit an affine dimer.

Proposition 3.2 Let $\Sigma = \{h_1, \dots, h_n\} \subseteq \mathbb{Z}^2$ be a multiset of primitive vectors and let P_Σ be the convex lattice polygon consisting of the pairwise antiparallel side segments $\pm h_1, \dots, \pm h_n$. Then P_Σ admits an affine dimer, i.e., P_Σ is the homology polygon of an admissible oriented line arrangement.

Additionally, the affine dimer may be taken to have $f_\times/f = 1/2$.

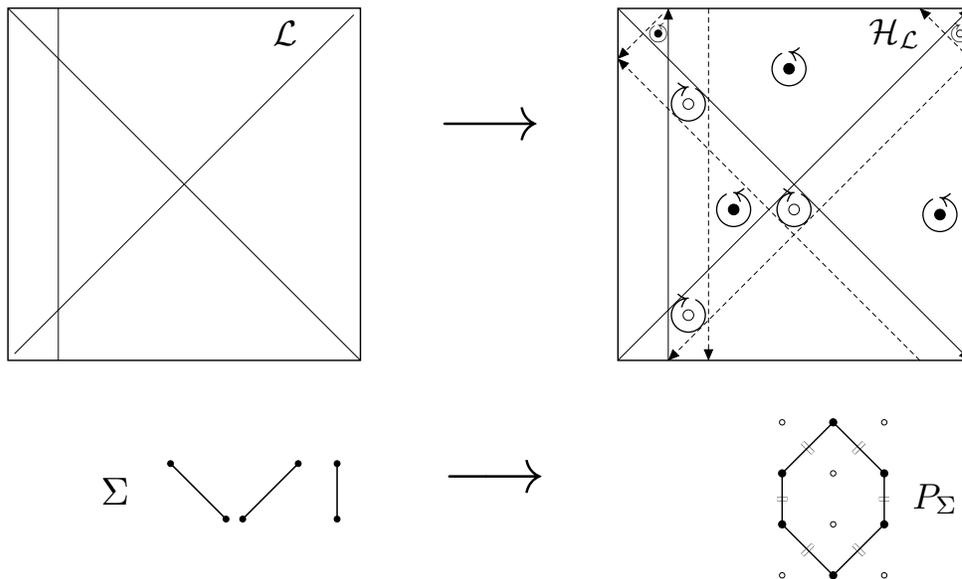


Figure 8: Illustration of the “double everything”-construction.

Proof. Let \mathcal{L} be any unoriented line arrangement in general position representing the homology classes Σ on \mathbb{T}^2 (up to sign). We construct an admissible oriented line arrangement $\mathcal{H}_\mathcal{L}$ as follows. First, add each line in \mathcal{L} to $\mathcal{H}_\mathcal{L}$. Then, for each $H \in \mathcal{L}$, add a line H^- to $\mathcal{H}_\mathcal{L}$ that is parallel and close enough to H , such that no lines intersect between H and H^- and $\mathcal{H}_\mathcal{L}$ is in general position.



Let V_\circ be the parallelograms corresponding to intersection points of \mathcal{L} and let V_\bullet be the faces corresponding to the original faces of \mathcal{L} . This gives an affine dimer $G = (V_\circ \sqcup V_\bullet, E)$ whose edges E encode the face-vertex incidence relations of \mathcal{L} (see Figure 8). By the discussion in Section 1.1 there is a choice of orientation for every line in $\mathcal{H}_\mathcal{L}$ such that the faces in V_\circ and V_\bullet are oriented clockwise and anticlockwise, respectively. This makes $\mathcal{H}_\mathcal{L}$ an admissible oriented line arrangement. Moreover, each pair (H, H^-) is oppositely oriented, so the homology polygon of $\mathcal{H}_\mathcal{L}$ is P_Σ , as required.

The inconsistently oriented faces of $\mathcal{H}_\mathcal{L}$ correspond to the line segments of \mathcal{L} and are all 4-gons. Thus, by Proposition 2.4, we have $f_\times/f = 1/2$. \square

3.3 Lifting

In this section we use column and row vectors for elements of a vector space and its dual, respectively.

Recall that $q : \mathbb{R}^2 \rightarrow \mathbb{T}^2 \cong \mathbb{R}^2/\mathbb{Z}^2$ is a universal cover of \mathbb{T}^2 . Let \mathcal{H} be an admissible oriented line arrangement on \mathbb{T}^2 . The preimage $q^{-1}(\mathcal{H})$ consists of all lifts of all the lines in \mathcal{H} . Moreover, each fundamental parallelogram on \mathbb{R}^2 spanned by $(1 \ 0)^T, (0 \ 1)^T$ contains exactly one representative copy of \mathcal{H} .

However, we may define a different fundamental parallelogram spanned by two elements of \mathbb{Z}^2 that gives a new universal cover of a torus on which it defines a new admissible oriented line arrangement. This is equivalent to first lifting \mathcal{H} to the universal cover \mathbb{R}^2 and then quotienting out by a general sublattice $\Lambda \leq \mathbb{Z}^2$. See Figure 9 for an example.

Let $\Lambda = \langle \alpha, \beta \rangle \leq \mathbb{Z}^2$ be a (non-degenerate) lattice and let \mathbb{T}_Λ^2 be the torus associated to the universal cover $q_\Lambda : \mathbb{R}^2 \rightarrow \mathbb{R}^2/\Lambda := \mathbb{T}_\Lambda^2$. Then $H_1(\mathbb{T}_\Lambda^2) \cong \mathbb{Z}\alpha \oplus \mathbb{Z}\beta$. We want to find the homology polygon of the admissible oriented line arrangement $\mathcal{H}_\Lambda := q_\Lambda(q^{-1}(\mathcal{H}))$ on \mathbb{T}_Λ^2 .

Note that since $\Lambda \leq \mathbb{Z}^2$, this construction gives a well-defined regular cover $q \circ q_\Lambda^{-1} : \mathbb{T}_\Lambda^2 \rightarrow \mathbb{T}^2$ of degree $|\text{covol}(\Lambda)|$, the volume of any fundamental parallelogram of Λ . This cover maps $q_\Lambda(q^{-1}(\mathcal{H}))$ onto \mathcal{H} so that \mathcal{H}_Λ is a regular cover of \mathcal{H} .



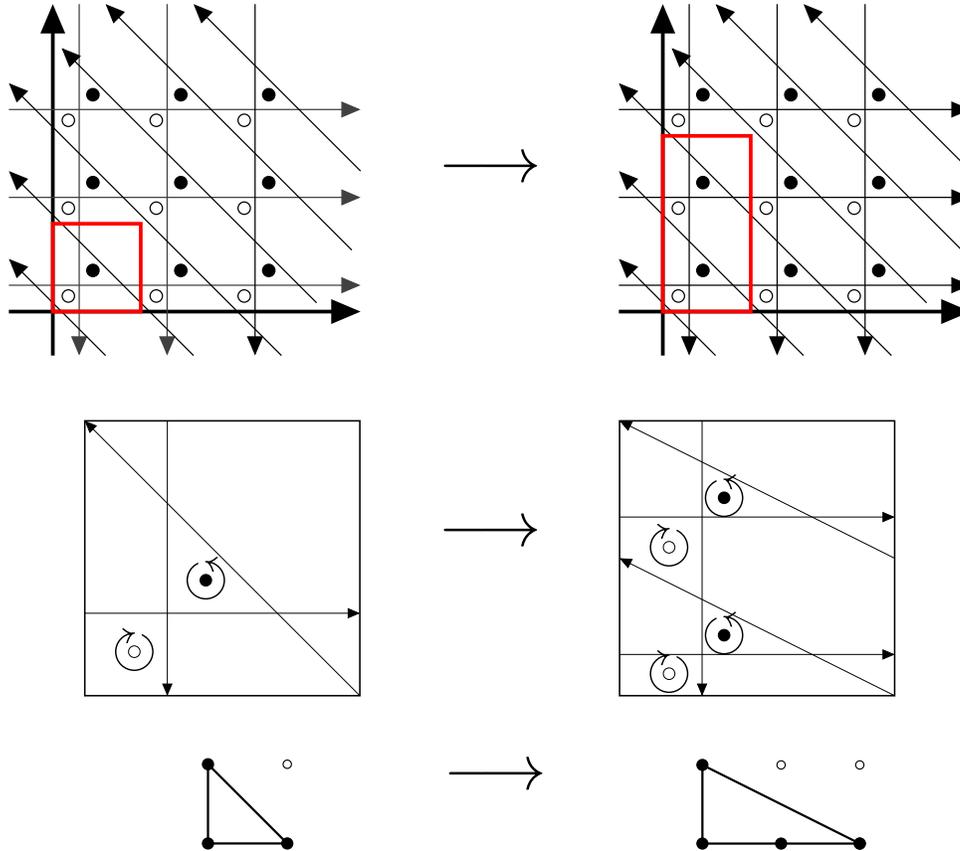


Figure 9: The lifting construction corresponding to the lattice $\Lambda = \langle (1 \ 0)^T, (0 \ 2)^T \rangle$. Top: change of fundamental parallelogram. Middle: from \mathbb{T}^2 to \mathbb{T}^2_Λ . Bottom: the new homology polygon.

Proposition 3.3 *Let P be the homology polygon of an admissible oriented line arrangement \mathcal{H} on \mathbb{T}^2 and let $\Lambda = \langle \alpha, \beta \rangle \leq \mathbb{Z}^2$ be a (non-degenerate) lattice with $\alpha = \begin{pmatrix} a \\ b \end{pmatrix}$ and $\beta = \begin{pmatrix} c \\ d \end{pmatrix}$. Let $A = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$ and let P_Λ be the homology polygon of \mathcal{H}_Λ on \mathbb{T}^2_Λ with respect to the basis $H_1(\mathbb{T}^2_\Lambda) \cong \mathbb{Z}\alpha \oplus \mathbb{Z}\beta$. Then*

$$P_\Lambda = \text{adj}(A)(P) = \begin{pmatrix} d & -c \\ -b & a \end{pmatrix} (P).$$

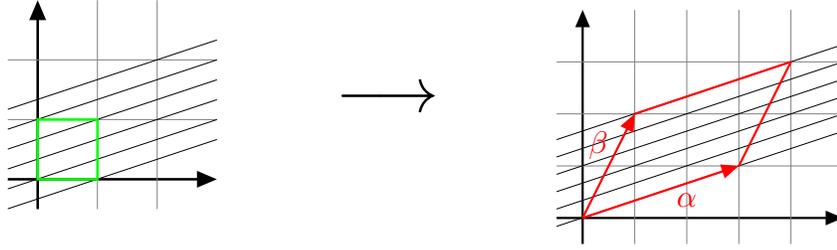


Figure 10: Illustration of the proof of Proposition 3.3 with $\alpha = (3 \ 1)^T$ and $\beta = (1 \ 2)^T$. In this case $\varphi((3 \ 1)^T) = (5 \ 0)^T$, confirming $\varphi(\alpha) = \det(A)\alpha \in H_1(\mathbb{T}_\Lambda^2) \cong \mathbb{Z}\alpha \oplus \mathbb{Z}\beta$.

Proof. Fix an orientation of \mathbb{T}^2 . For two transversal loops γ_1 and γ_2 on a torus let $\iota(\gamma_1, \gamma_2)$ be their signed intersection number, which is invariant under homotopy. By Poincaré duality we have an isomorphism

$$\begin{aligned} i : H_1(\mathbb{T}^2) &\xrightarrow{\cong} H^1(\mathbb{T}^2) \\ [\gamma] &\longmapsto \iota([\gamma], \cdot) \end{aligned}$$

and similarly $i_\Lambda : H_1(\mathbb{T}_\Lambda^2) \xrightarrow{\cong} H^1(\mathbb{T}_\Lambda^2)$. As discussed above, the map $\pi_\Lambda := q \circ q_\Lambda^{-1} : \mathbb{T}_\Lambda^2 \rightarrow \mathbb{T}^2$ is a regular cover, restricting to a regular cover \mathcal{H}_Λ of \mathcal{H} . Let $\varphi := i_\Lambda^{-1} \circ \pi_\Lambda^* \circ i$.

$$\begin{array}{ccc} H_1(\mathbb{T}^2) & \xrightarrow{\varphi} & H_1(\mathbb{T}_\Lambda^2) \\ \cong \downarrow i & & \cong \downarrow i_\Lambda \\ H^1(\mathbb{T}^2) & \xrightarrow{\pi_\Lambda^*} & H^1(\mathbb{T}_\Lambda^2) \end{array} \quad (1)$$

We shall show that $\varphi([l]) = [\pi_\Lambda^{-1}(l)]$ for every $l \in Z_1(\mathbb{T}^2)$. By Poincaré duality this is equivalent to

$$\iota(l, \pi_\Lambda(g)) = \iota(\pi_\Lambda^{-1}(l), g) \quad (2)$$

for all $l \in Z_1(\mathbb{T}^2)$ and $g \in Z_1(\mathbb{T}_\Lambda^2)$. Indeed, π_Λ is orientation preserving and for $l : S^1 \rightarrow \mathbb{T}^2$ and $g : S^1 \rightarrow \mathbb{T}_\Lambda^2$ we have $l(s) = \pi_\Lambda(g(t))$ if and only if $g(t) = \hat{l}(s)$ for some lift \hat{l} of l along π_Λ , proving (2).

It remains to show that φ in (1) is given by the matrix $\text{adj}(A)$. Working in the basis of $H^1(\mathbb{Z}_\Lambda^2)$ dual to $H_1(\mathbb{Z}_\Lambda^2) \cong \mathbb{Z}\alpha \oplus \mathbb{Z}\beta$ we obtain the desired result:

$$\begin{bmatrix} i \begin{pmatrix} 1 \\ 0 \end{pmatrix} = (0 \ 1) \\ i \begin{pmatrix} 0 \\ 1 \end{pmatrix} = (-1 \ 0) \end{bmatrix} \implies \begin{bmatrix} \pi_\Lambda^* \circ i \begin{pmatrix} 1 \\ 0 \end{pmatrix} = (b \ d) \\ \pi_\Lambda^* \circ i \begin{pmatrix} 0 \\ 1 \end{pmatrix} = (-a \ -c) \end{bmatrix} \implies \begin{bmatrix} \varphi \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} d \\ -b \end{pmatrix} \\ \varphi \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -c \\ a \end{pmatrix} \end{bmatrix}$$

□

The map $A \mapsto \text{adj}(A)$ on $\{A \in \mathbb{Z}^{2 \times 2} : \det(A) \neq 0\}$ is surjective. Thus:



Corollary 3.4 *If P is the homology polygon of an affine dimer and $B \in \mathbb{Z}^{2 \times 2}$ with $\det(B) \neq 0$, then $B(P)$ is also the homology polygon of an affine dimer.*

We already knew this for $B \in GL_2(\mathbb{Z})$ because $GL_2(\mathbb{Z})$ acts by linear automorphisms on \mathbb{T}^2 . This corollary is a generalisation.

4 Affine Dimer Search Algorithm

The class of homology polygons obtained from the constructions in Section 3 is not too big. Indeed, if P is a homology polygon obtained from Proposition 3.1 or Proposition 3.2 then P has a pair of antiparallel side segments. If P is obtained by lifting using a matrix $B \in \mathbb{Z}^{2 \times 2}$ with $\det(B) \neq 0$ as in Corollary 3.4, and if the non-primitive side segments of P are $p_1, \dots, p_m \in \mathbb{Z}^2$, then $\det(B) \mid \det(p_i, p_j)$ for all i, j . For this construction to deliver a new $GL_2(\mathbb{Z})$ equivalence class of convex lattice polygons, we require $\det(B) \neq \{0, \pm 1\}$. Thus, the integers $\det(p_i, p_j)$ all have a common prime factor, which is a rare trait for a convex lattice polygon P .

Therefore, we developed a computer program with GUI to manipulate line arrangements on the torus and check whether a given convex lattice polygon admits an affine dimer. This section summarises the algorithms used.

The programming was done primarily in Java, using the library `JGraphT` [10] for standard graph algorithms and `polymake` [13] to work with cell decompositions of \mathbb{R}^n .

4.1 Checking a single arrangement

Given a line arrangement $\mathcal{H} = \{H_1, \dots, H_n\}$ in general position on \mathbb{T}^2 with $\sum_{i=1}^n [H_i] = 0$, the following algorithm determines whether it corresponds to an affine dimer.

1. Calculate all intersection points. For a pair (H_1, H_2) of lines, this is done by setting $[H_1] = (1, 0)$ via an action of $SL_2(\mathbb{Z})$. This simplified configuration is dealt with by inspection. The number of intersection points of H_1 and H_2 is $|\det([H_1], [H_2])|$.
2. For each line $H \in \mathcal{H}$, determine the order of the intersection points on H . Again, this is done by first setting $[H_1] = (1, 0)$ via $SL_2(\mathbb{Z})$. Thus, we obtain the side segment data of \mathcal{H} .
3. For each intersection point of each side segment, determine the next side segment at that point in clockwise and anticlockwise order. Thus, we obtain the face data of \mathcal{H} . Abstract this to a graph structure in which two faces are neighbours if and only if they share a vertex.
4. Determine the number k of bipartite connected components of the obtained graph.
5. The obtained graph has exactly two connected components, which can be seen by considering intersection numbers modulo 2. Hence, there are three cases:



- If $k = 2$ then the homology polygon of \mathcal{H} is a parallelogram by Lemma 4.1 below, which is already known to admit an affine dimer by the “double everything”-construction of Proposition 3.2.
- If $k = 1$ then there is a choice of orientation for each line in \mathcal{H} making the arrangement admissible. A further check unveils whether this is compatible with the given orientations. If not, a dimer for a different homology polygon has been found. See Figure 12 for an example of why this is necessary.
- If $k = 0$ then the configuration is not admissible and this cannot be fixed by re-orienting the lines.

This algorithm has linear time and space complexity $\mathcal{O}(f)$ by Proposition 2.2 (iv), where f is the number of faces of \mathcal{H} .

Lemma 4.1 *If $k = 2$ in the above algorithm, then the homology polygon P of \mathcal{H} is a parallelogram.*

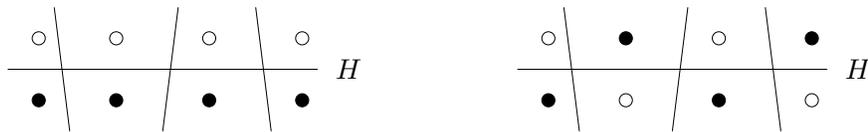


Figure 11: The two types of lines in the case $k = 2$.

Proof. Since the obtained graph is bipartite and $k = 2$, each line $H \in \mathcal{H}$ is of one of the two types depicted in Figure 11. Thus, no two lines of the same type intersect, so all lines of the same type are parallel (or antiparallel). Thus, there are at most 4 homology classes and so P is a parallelogram. \square

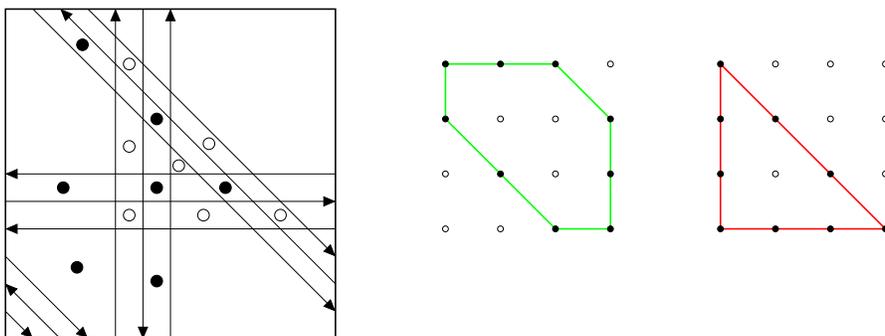


Figure 12: The admissible arrangement to the left appears in the moduli space of the red (right) homology polygon and has $k = 1$ in the above algorithm. However, it represents the green (middle) homology polygon.

4.2 Moduli space of line arrangements

Checking whether a convex lattice polygon P admits an affine dimer is more difficult as there are infinitely many line arrangements in general position realising P . However, only finitely many of them represent different combinatorial configurations. We consider two arrangements to be combinatorially the same if one of them can be obtained from the other by continuously translating some lines without ever creating a triple intersection point or two coinciding parallels. This notion is formalised by the moduli space \mathcal{M} of line arrangements.

Lemma 4.2 *Let $\alpha \in \mathbb{Z}^n$, $c \in \mathbb{R}$, and $\hat{H} = \{x \in \mathbb{R}^n : \langle x, \alpha \rangle = c\}$ be a hyperplane. Let $q : \mathbb{R}^n \rightarrow \mathbb{T}^n$ be the quotient map and $H = q(\hat{H})$. Then $x + \mathbb{Z}^n \in H$ if and only if $\langle x, \alpha \rangle \in c + \gcd(\alpha)\mathbb{Z}$.*

Proof. \mathbb{Z} is a Euclidean domain, so the ideal equation $(\alpha_1, \dots, \alpha_n) = \gcd(\alpha)\mathbb{Z}$ holds. \square

Therefore, given a primitive homology class $\alpha \in \mathbb{Z}^2$, the set of lines realising this homology class is parametrised uniquely by $c \in \mathbb{R}/\mathbb{Z} \cong \mathbb{T}$, since $\gcd(\alpha) = 1$. Therefore:

Definition 4.3 *The moduli space \mathcal{M} of line arrangements on \mathbb{T}^2 consisting of n lines with prescribed homology class is topologically \mathbb{T}^n . More precisely, if the primitive side segments of P are $h_1, \dots, h_n \in \mathbb{Z}^2$ and α_i is the clockwise rotation of h_i by $\pi/2$ then the correspondence is*

$$(c_1, \dots, c_n) + \mathbb{Z}^n \in \mathbb{T}^n \cong \mathcal{M} \quad \longleftrightarrow \quad \left[\begin{array}{l} \mathcal{H} = \{H_1, \dots, H_n\}, \quad [H_i] = h_i, \\ H_i = \{x + \mathbb{Z}^2 : \langle x, \alpha_i \rangle \in c_i + \mathbb{Z}\} \end{array} \right].$$

Let $C \subset \mathcal{M}$ be the locus where \mathcal{H} is not in general position. This happens either when three or more lines intersect in a point or when two (anti)parallel lines have the same parameter.

For each pair $\{h_i, h_j\}$ with $i \neq j$ and $h_i \parallel h_j$, H_i and H_j coincide if and only if $c_i = c_j$ on \mathbb{T} . This happens if and only if (c_1, \dots, c_n) lies on the hyperplane $C_{i,j} : X_i - X_j = 0$ on \mathcal{M} .

For each triple $\{h_i, h_j, h_k\}$ of pairwise non-(anti)parallel homology classes, the lines H_i, H_j, H_k intersect in a common point $x \in \mathbb{T}^2$ if and only if

$$\langle x, \alpha_i \rangle \equiv c_i, \quad \langle x, \alpha_j \rangle \equiv c_j, \quad \text{and} \quad \langle x, \alpha_k \rangle \equiv c_k \pmod{\mathbb{Z}}.$$

This holds for some x if and only if

$$(c_i, c_j, c_k) \in \text{Im}(A) + \mathbb{Z}^3 \text{ for the } 3 \times 2 \text{ integer matrix } A = \begin{pmatrix} \leftarrow \alpha_i \rightarrow \\ \leftarrow \alpha_j \rightarrow \\ \leftarrow \alpha_k \rightarrow \end{pmatrix} =: \begin{pmatrix} \uparrow & \uparrow \\ e_{ijk} & f_{ijk} \\ \downarrow & \downarrow \end{pmatrix},$$

or equivalently $\langle (c_i, c_j, c_k), e_{ijk} \wedge f_{ijk} \rangle \in \gcd(e_{ijk} \wedge f_{ijk})\mathbb{Z}$ by Lemma 4.2. This defines a hyperplane $D_{i,j,k} \subset \mathbb{T}^n \cong \mathcal{M}$. Thus, the locus $C \subset \mathcal{M}$ of degenerate arrangements is



given by the hyperplane arrangement

$$C := \left(\bigcup_{\substack{h_i \parallel h_j \\ i \neq j}} C_{i,j} \right) \cup \left(\bigcup_{\substack{h_i, h_j, h_k \text{ pairwise} \\ \text{non-(anti)parallel}}} D_{i,j,k} \right) \subset \mathcal{M}.$$

To determine if P admits an affine dimer, one therefore needs to apply the algorithm of the previous subsection to one arrangement (c_1, \dots, c_n) of each connected component of the complement $\mathcal{M} \setminus C$. It remains to enumerate the connected components of $\mathcal{M} \setminus C$.

Note also that the coordinates of $e_{ijk} \wedge f_{ijk}$ in the definition of $D_{i,j,k}$ are of the form $\det(h_{i_1}, h_{i_2})$. Thus, the structure $C \subset \mathcal{M}$ only depends on the $GL_2(\mathbb{Z})$ equivalence class the homology polygon P .

4.2.1 Cell decomposition approach

One way of enumerating the components of $\mathcal{M} \setminus C$ is as follows.

1. Lift each constituting hyperplane $H = D_{i,j,k}$ or $C_{i,j}$ of C to several hyperplanes $\{H_1, \dots, H_l\}$ in \mathbb{R}^n such that $H = [0, 1]^n \cap \{H_1, \dots, H_l\} / \mathbb{Z}^n$. (For $H = C_{i,j}$ we have $l = 1$. For $H = D_{i,j,k}$ with normal vector $\nu := e_{ijk} \wedge f_{ijk}$ we have $l \leq 1 + \sqrt{3} \|\nu\| / \gcd(\nu)$.) Thus, we obtain a (finite) hyperplane arrangement $\hat{C} \subseteq \mathbb{R}^n$ such that $q(\hat{C}) = C$.
2. Use a cell decomposition algorithm to find the connected components of $\mathbb{R}^n \setminus \hat{C}$.
3. Pick one point in each cell and check if the corresponding configuration is admissible.

We implemented this using `polymake` [13], but the resulting algorithm was not efficient enough to deliver results for some lattice polygons P of interest (such as the conjectured counterexamples of Forsgård for $k = 3, 4$ ([5], Section 4)). This is no surprise since C consists of $\mathcal{O}(n^3)$ hyperplanes, each lifting to an arbitrarily large number of hyperplanes in \hat{C} . For a hyperplane arrangement in general position on \mathbb{R}^n the number of components of the complement is a degree n polynomial in the number of hyperplanes. Thus, the number of cells could be $\mathcal{O}(n^{3n})$ or even higher, depending on the number of lifts when constructing \hat{C} from C .

It might be possible to optimize this using properties of \mathbb{T}^n or the fact that we are only working on $[0, 1]^n$.

4.2.2 Mesh approach

The cell decomposition algorithm of `polymake` gives us a proper cell decomposition, whereas we only really need one point in each cell of $\mathcal{M} \setminus C$. There is a smallest constant $m(C) \in \mathbb{N}_{>0}$ such that every component of $\mathcal{M} \setminus C$ contains a point of $q(m(C)^{-1}\mathbb{Z}^n)$. Thus, we are able to finish by checking $\mathcal{O}(m(C)^n)$ configurations and without calculating any cell decompositions.



However, even when C just consists of two lines on \mathbb{T}^2 , there are configurations of arbitrarily large $m(C)$. E.g., take lines of homology class $(1, 0)$ and $(1, l)$ with $l \rightarrow \infty$. This configuration has l faces, so $m(C) \geq l$. Note also that this cannot be cured by applying a smart choice of $A \in GL_2(\mathbb{Z})$ to C since the number of faces is invariant.

4.2.3 Reducing dimension by two

In both approaches described above, we may reduce the dimension by two by restricting our attention to the subtorus $\mathbb{T}^{n-2} \subset \mathcal{M}$ with $c_1 \equiv c_2 \equiv 0 \pmod{\mathbb{Z}}$. This corresponds to translating all arrangements so that H_1, H_2 intersect at the bottom left corner of the fundamental parallelogram, where $h_1 \nparallel h_2$ without loss of generality.

4.3 Randomized search & admissible volume

We may choose random vectors $(c_1, \dots, c_n) \in \mathcal{M}$ and check each configuration until we find an affine dimer, or stop after a certain number of trials. This approach led to the discovery of many non-trivial affine dimers of genus 1 and 2 presented in Section 5.

Definition 4.4 *Let $\mathcal{A} \subset \mathcal{M} \setminus C$ be the locus of admissible oriented line arrangements. The admissible volume of P is $\text{vol}(\mathcal{A})$.*

Note that $\text{vol}(\mathcal{A})$ only depends on the $GL_2(\mathbb{Z})$ equivalence class of the homology polygon P , since C is $GL_2(\mathbb{Z})$ invariant. Thus, it makes sense to talk about the admissible volume of an $GL_2(\mathbb{Z})$ equivalence class of convex lattice polygons.

Using this randomized approach it is possible to estimate the admissible volume of P . For some homology polygons in Section 5 we had $\text{vol}(\mathcal{A}) < 0.01$, making it highly unlikely that we could have found their affine dimers by hand, and justifying our computational approach.

Furthermore, if $\mathbb{T}^{n-2} \subset \mathcal{M}$ denotes the subtorus with $c_1 \equiv c_2 \equiv 0 \pmod{\mathbb{Z}}$, then $\text{vol}_{\mathbb{T}^n}(\mathcal{A}) = \text{vol}_{\mathbb{T}^{n-2}}(\mathcal{A} \cap \mathbb{T}^{n-2})$. This is because the $(n-2)$ -dimensional fibers are isomorphic via global translation of the arrangements, and translation is an isometry on flat tori. Thus, we may speed up the estimation of $\text{vol}(\mathcal{A})$ (and thus our search for affine dimers) by reducing the dimension by two.

Finding bounds for $\text{vol}(\mathcal{A})$ in terms of P might allow us to answer Question 1.10 for bigger classes of polygons. For example, we have the following result for parallelograms, which might be generalised in future.

Proposition 4.5 *Let P be an $a \times b$ lattice parallelogram with $n = 2a + 2b$ primitive side segments. Then*

$$\text{vol}(\mathcal{A}) = 4 \binom{2a}{a}^{-1} \binom{2b}{b}^{-1}.$$

Proof. We say that P has $2a$ horizontal and $2b$ vertical side segments, a (respectively b) of each orientation. An arrangement $(c_1, \dots, c_n) \in \mathcal{M} \setminus C$ is admissible if and only if both the vertical lines and the horizontal lines have alternating orientations. There are $\binom{2a}{a}$ orders



of orientations for the horizontal lines, exactly two of which are alternating (and similarly for the vertical lines). The result follows since all orderings of horizontal (respectively vertical) lines are equally likely and from independence of vertical and horizontal lines. \square

5 Triangles and Genus ≤ 2

We now apply the constructions of Section 3 and the algorithms of Section 4 to exhibit various families of convex lattice polygons that admit an affine dimer. We also record some estimates of their admissible volumes (see Section 4.3), indicating how hard it would be to find their affine dimers by hand without our constructions and algorithms.

5.1 Triangles

We begin with an application of the lifting construction in Section 3.3.

Proposition 5.1 *Let P be a lattice triangle (possibly with more than three primitive side segments). Then P admits an affine dimer. Moreover, it admits an affine dimer that is lifted from the one in Figure 6.*

Proof. Let the triangle P be spanned by the vectors $(a, b), (c, d) \in \mathbb{Z}^2$, not necessarily primitive as we allow more than three primitive side segments. Then an admissible oriented line configuration with homology polygon P is obtained by applying the matrix $B = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$ to the triangle spanned by $(1, 0)$ and $(0, 1)$ using Corollary 3.4. \square

5.2 Genus Zero

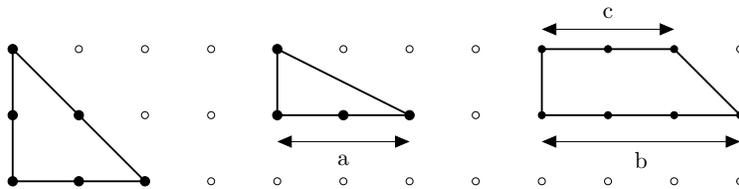


Figure 13: The three families of equivalence classes of convex lattice polygons with no interior lattice points, where a, b, c are positive integers [18].

Proposition 5.2 *Let P be a convex lattice polygon with no interior lattice points. Then P admits an affine dimer.*

Proof. By [18] the equivalence classes of convex lattice polygons with no interior lattice points are those displayed in Figure 13. By Proposition 5.1, the triangles all admit an affine dimer.

For the trapezoid given by $b, c \in \mathbb{Z}_{>0}$, there are two cases. If $b = c$ then P consists of pairwise antiparallel edges, so admits an affine dimer by Proposition 3.2. If $b \neq c$ and



without loss of generality $b > c$ then P is obtained from the triangle in Figure 13 with $a = c$ by adding $b - c$ pairs of antiparallel edges parallel to $(1, 0)$. Thus, P admits an affine dimer by Proposition 5.1 and Proposition 3.1. \square

Even for the genus 0 triangles the admissible volume decays quickly:

Proposition 5.3 *Let P be the $a \times 1$ triangle in Figure 13. The admissible volume of $\mathcal{A} \subset \mathcal{M} \setminus C$ is*

$$\text{vol}(\mathcal{A}) = a!/a^a.$$

Proof. Let H_1, \dots, H_a be the lines of homology class $(1, 0)$ and let V and S be the lines of homology class $(0, -1)$ and $(-a, 1)$, respectively. Then S subdivides V into a segments s_1, \dots, s_a of equal length. Since the arrangement is in general position, there is a function $f : \{1, \dots, a\} \rightarrow \{1, \dots, a\}$ such that H_i intersects V in the segment $s_{f(i)}$. By inspection, the arrangement is admissible if and only if f is injective. Since every f is equally likely for $(c_1, \dots, c_{a+2}) \in \mathcal{M} \setminus C$ chosen uniformly at random, we have

$$\text{vol}(\mathcal{A}) = \mathbb{P}(f \text{ is injective}) = a!/a^a.$$

\square

5.3 Genus One

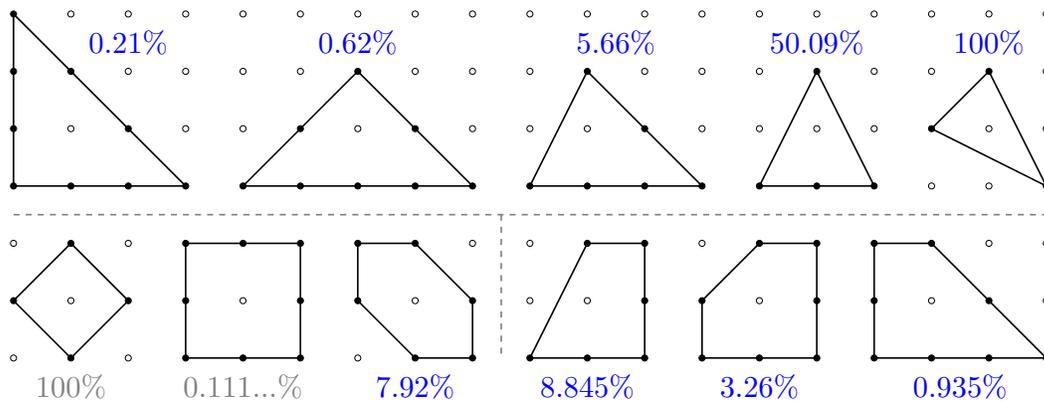


Figure 14: Equivalence class representatives of convex lattice polygons with one interior point that are triangles (top), consist of pairwise antiparallel side segments (bottom left), or have a pair of antiparallel side segments which are parallel to at least one other side segment (bottom right). The blue numbers are estimates ($\geq 2 \cdot 10^4$ trials) of the admissible volume of the polygon, indicating how hard it would be to find affine dimers by hand. The two bottom left numbers are exact by Proposition 4.5.

Proposition 5.4 *Let P be a convex lattice polygon with exactly one interior lattice point. Then P admits an affine dimer.*



Proof. There are 16 equivalence classes of convex lattice polygons with exactly one interior point [17], [18]. As seen in Figure 14, four of them are triangles, which admit an affine dimer by Proposition 5.1. Three of them consist of pairwise antiparallel side segments, so admit an affine dimer by Proposition 3.2. Another three of them are obtained by adding a pair of antiparallel side segments $(\pm 1, 0)$ to convex lattice polygons containing $(1, 0)$ as a side segment which are already known to be dimers since they have no interior points. These admit an affine dimer by Proposition 3.1.

There are five equivalence classes left whose affine dimers are displayed in Figure 15 & 16. □

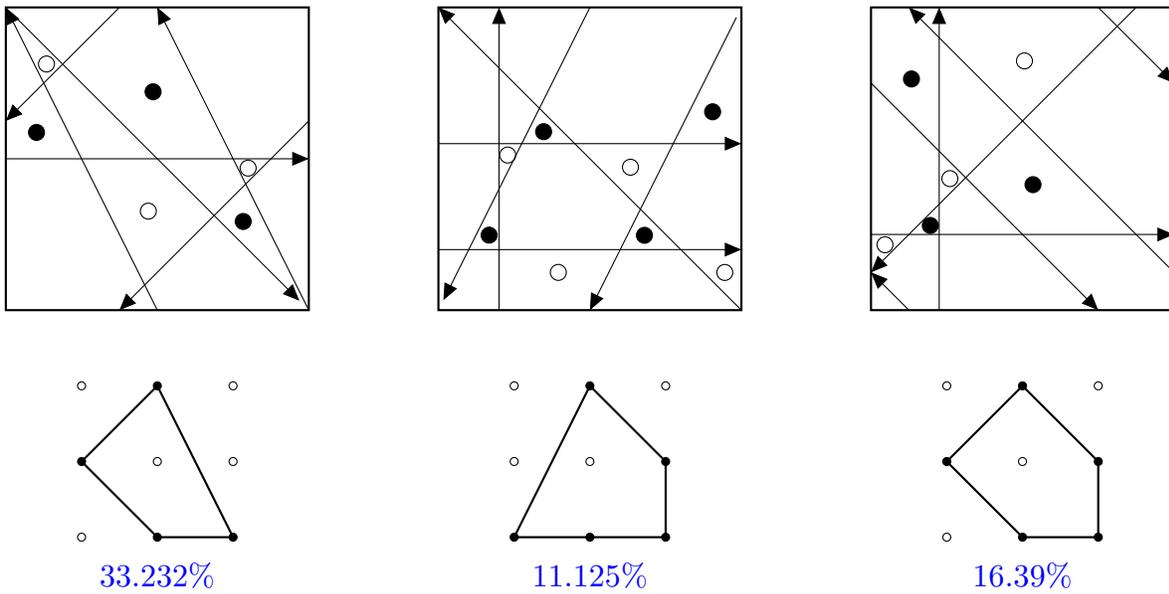


Figure 15: The five remaining equivalence classes of convex lattice polygons with one interior lattice point and affine dimers for them (continued in Figure 16). The blue numbers are estimates ($\geq 2 \cdot 10^4$ trials) of the admissible volume.

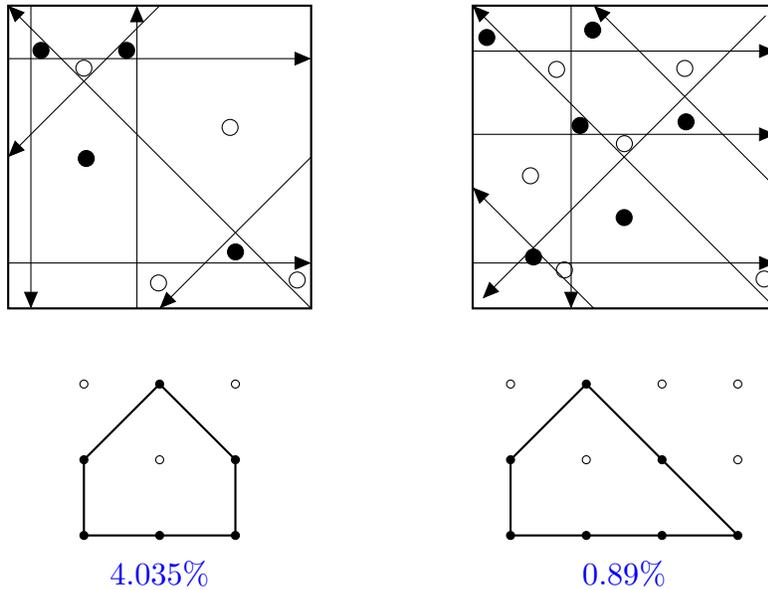


Figure 16: Continuation of Figure 15.

5.4 Genus Two

The same holds for two interior points.

Proposition 5.5 *Let P be a convex lattice polygon with exactly two interior lattice points. Then P admits an affine dimer.*

Proof. A classification up to equivalence of convex lattice polygons with two interior lattice points is provided by [21].

There are five classes of triangles, all of which admit an affine dimer by Proposition 5.1.

There are 19 classes of quadrilaterals. Three of them are parallelograms and thus admit an affine dimer by Proposition 3.1. Six of them are obtained by adding a pair of antiparallel edges parallel to an existing edge to a convex lattice polygon of genus 0 or 1, and therefore admit an affine dimer by Propositions 5.2, 5.4, and 3.1. The other 10 classes of quadrilaterals were checked manually to admit an affine dimer (see below).

Similarly, all sixteen classes of pentagons and five classes of hexagons admit an affine dimer (see Figure 17 for an example). There are no convex lattice n -gons with two interior lattice points and $n > 6$. The 19 classes that required computer-aided verification can be found online at <https://jeffhicks.net/files/DHolmesSupplemental.pdf>. \square

This completes the proof of Theorem B.



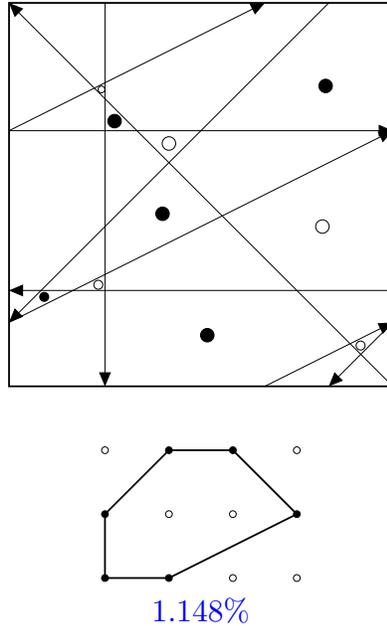


Figure 17: The only genus 2 hexagon whose dimer cannot be constructed from a lower genus dimer using the constructions of Section 3. The blue number is an estimate (10^5 trials) of the admissible volume.

Acknowledgments

The author would like to thank the *London Mathematical Society* and the *Department of Pure Mathematics and Mathematical Statistics, University of Cambridge* for their financial support in the form of an Undergraduate Research Bursary, and Jeff Hicks for suggesting and mentoring this project, and for providing invaluable feedback and many fruitful discussions. Finally, we would like to thank the referees for their helpful feedback, contributing multiple improvements to this paper.

References

- [1] M. Aigner, G.M. Ziegler, *Proofs from THE BOOK* (5th ed.), Springer, 2014.
- [2] T.F.N. Chan, Dimer models and their characteristic polygons, 2016, available online at the URL: https://web.maths.unsw.edu.au/~danielch/thesis/Chan_timothy.pdf
- [3] D. Cimasoni, The geometry of dimer models, *Winter Braids Lect. Notes*, **1** (2014), talk no. 2.
- [4] R. Diaz, S. Robins, Pick's Formula via the Weierstrass \wp -Function, *Amer. Math. Monthly*, **102** (1995), 431–437.
- [5] J. Forsgård, On dimer models and coamoebas, *Ann. Inst. Henri Poincaré D*, **6** (2019), 199–219.
- [6] J. Forsgård, P. Johansson, Coamoebas and line arrangements in dimension two, *Math. Z.*, **278** (2014), 25–38.
- [7] M. Futaki and K. Ueda, Dimer models and homological mirror symmetry for triangles, 2010, available online at the URL: <https://arxiv.org/abs/1004.3620>



- [8] D.R. Gulotta, Properly ordered dimers, R-charges, and an efficient inverse algorithm, *J. High Energy Phys.*, **2008** (2008), article 14.
- [9] J. Hicks, Tropical Lagrangians in toric del-Pezzo surfaces, *Sel. Math. New Ser.*, **27** (2021), article 3.
- [10] JGraphT. A Java library of graph theory data structures and algorithms, available online at the URL: <https://jgrapht.org/>
- [11] P. Johansson, The argument cycle and the coamoeba, *Complex Var. Elliptic Equ.*, **58** (2013), 373–384.
- [12] P.W. Kasteleyn, The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice, *Physica*, **27** (1961), 1209–1225.
- [13] L. Kastner, M. Panizzut, Hyperplane Arrangements in polymake, in *Math. Software – ICMS 2020*, Springer, 2020, 232–240.
- [14] R. Kenyon, A. Okounkov, S. Sheffield, Dimers and amoebae, *Ann. of Math. (2)*, **163** (2006), 1019–1056.
- [15] G. Mikhalkin, Enumerative Tropical Algebraic Geometry in \mathbb{R}^2 , *J. Amer. Math. Soc.*, **18** (2005), 313–377.
- [16] A. Okounkov, N. Reshetikhin, C. Vafa, Quantum Calabi-Yau and classical crystals, *Progr. Math.*, **244** (2006), 597–618.
- [17] B. Poonen, F. Rodriguez-Villegas, Lattice Polygons and the Number 12, *Amer. Math. Monthly*, **107** (2000), 238–250.
- [18] S. Rabinowitz, A census of convex lattice polygons with at most one interior lattice point, *Ars Combin.*, **28** (1989), 83–96.
- [19] K. Ueda, M. Yamazaki, A Note on Dimer Models and McKay Quivers, *Commun. Math. Phys.*, **301** (2011), 723–747.
- [20] K. Ueda, M. Yamazaki, Homological mirror symmetry for toric orbifolds of toric del Pezzo surfaces, *J. Reine Angew. Math.*, **680** (2013), 1–22.
- [21] X. Wei, R. Ding, Lattice polygons with two interior lattice points, *Math. Notes*, **91** (2012), 868–877.

Daniel Holmes

DPMMS, University of Cambridge
 Wilberforce Road
 Cambridge, CB3 0WB, United Kingdom
 E-mail: dh604@cam.ac.uk

Received: October 7, 2021 **Accepted:** January 5, 2022
Communicated by Matthias Beck

