TOPOLOGICAL INDUCTIVE CONSTRUCTIONS FOR TIGHT SURFACE GRAPHS

JAMES CRUICKSHANK, DEREK KITSON, STEPHEN C. POWER, AND QAYS SHAKIR

Abstract. We investigate properties of sparse and tight surface graphs. In particular we derive topological inductive constructions for $(2,2)$-tight surface graphs in the case of the sphere, the plane, the twice punctured sphere and the torus. In the case of the torus we identify all 116 irreducible base graphs and provide a geometric application involving contact graphs of configurations of circular arcs.

1. Introduction

Inductive characterisations of various families of graphs play an important role in many parts of graph theory. A graph G is said to be $(2,2)$-sparse if for any nonempty $V' \subset V(G)$, we have $|E(V')| \leq 2|V'| - 2$. If, in addition, $|E(G)| = 2|V(G)| - 2$ we say that G is $(2,2)$-tight. Such graphs arise naturally in various parts of geometric graph theory, including framework rigidity, circle packings, and also in graph drawing.

We will derive inductive characterisations of $(2,2)$-tight graphs that are embedded without edge crossings in certain orientable surfaces of genus at most 1. Our characterisations will be based on edge contractions and are in the spirit of well known results of Barnette, Nakamoto and others ([2, 16, 17]) on irreducible triangulations and quadrangulations of various surfaces. It may be worth noting, for example, that a plane simple graph is a quadrangulation of the plane if and only if it is $(2,4)$-tight: that is, $|E(G)| = 2|V(G)| - 4$ and for all $V' \subset V(G)$ such that $|V'| \geq 3$ we have $|E(V')| \leq 2|V'| - 4$. Thus it is clear that our results are related to, but distinct from, existing results on quadrangulations.

Since our graphs are embedded, we consider inductive characterisations based on topological edge contractions - that is to say that the contraction preserves the embedding of the graph (precise definitions given below). This is a key point, since the well-known inductive characterisation of simple $(2,2)$-tight graphs by Nixon, Owen and Power is purely graph theoretic ([18]). To further illustrate the significance of this we give an application of our main result to a recognition problem in graph drawing. The topological nature of the inductive characterisation is crucial in this context.

Finally we note that there are related topological inductive characterisations of Laman graphs in the literature already which have interesting geometric applications to pseudotriangulations and auxetic structures ([7, 11]).

2010 *Mathematics Subject Classification.* 05C10, 05C62, 52C30.

Key words and phrases. graph, surface, torus graph, rotation system, sparse graph, tight graph, vertex splitting, inductive construction, contact graph, contacts of circular arcs.

The second and third authors were supported by the Engineering and Physical Sciences Research Council [grant number EP/P01108X/1].

The fourth author gratefully acknowledges the financial support from the Iraqi Ministry of Higher Education and Scientific Research and Middle Technical University, Baghdad.

1.1. **Summary of main results.** Section 2 and the first part of Section 3 are background material for the rest of the paper. The main contributions of the paper are as follows:

- Theorem 3.4 presents an elementary but very useful principle concerning sparsity counts and graphs embedded in surfaces. While related results and special cases already exist in the literature, our statement and proof emphasises that this is a general principle that applies to a wide range of sparsity counts and to surfaces of all genus.
- In Section 4 we analyse the quadrilateral contraction move. This operation is well known in the context of quadrangulations. Here we examine its properties with respect to $(2,2)$-sparsity and prove some structural results about non contractible quadrilaterals in this context.
- Theorem 6.6 shows that if G is an irreducible $(2,2)$-tight surface graph, then any $(2,2)$-tight subgraph is also irreducible. This holds for surfaces of arbitrary genus.
- We give topological inductive characterisations of $(2,2)$-tight graphs embedded in the sphere, plane, annulus (Theorem 5.4) and the torus (Theorem 7.6). The first two of these are relatively standard, whereas the latter two characterisations are new. In the case of Theorem 7.6 we have also identified the 116 irreducible $(2,2)$-tight torus graphs. We do not give an explicit description of these graphs in the paper for reasons of space, but the reader is referred to [5] and [20] for details.
- Finally we present an application of our results to a recognition problem in geometric graph theory. Specifically we show that every $(2,2)$-tight torus graph can be realised as the contact graph of a collection of non overlapping circular arcs in the flat torus. We note that by passing to the universal cover this result may also be interpreted as a recognition result for contact graphs of doubly periodic collections of circular arcs in the plane.

We also conjecture a generalisation of our main results (Conjecture 5.1) to the set of irreducible $(2,2)$-tight surface graphs for surfaces of arbitrary genus. This would be analogous to results of Barnette, Nakamoto and others on triangulations and quadrangulations of surfaces. As noted above and at appropriate points in the paper, many of our results are valid for a range of sparsity counts and for surfaces of arbitrary genus and these will be useful in future investigations of this conjecture.

2. Graphs, surfaces and embeddings

In this section we fix our conventions and terminology regarding topological graphs. Throughout, we use the word graph for a finite undirected multigraph. So loop and parallel edges are allowed a priori, although loop edges will not arise in most of the cases of interest. If Γ is a graph then $|\Gamma|$ is its geometric realisation.

Suppose that Σ is a real orientable 2-dimensional manifold without boundary. A Σ-graph, or surface graph, is a pair (Γ, φ) where Γ is a graph and $\varphi : |\Gamma| \to \Sigma$ is a continuous embedding of the geometric realisation of Γ in Σ. Given Σ_i-graphs (Γ_i, φ_i) for $i = 1, 2$ we say that they are isomorphic if there is a homeomorphism $h : \Sigma_1 \to \Sigma_2$ and a graph isomorphism $g : \Gamma_1 \to \Gamma_2$ such that $h \circ \varphi_1 = \varphi_2 \circ |g|$, where $|g|$ is the induced homeomorphism $|\Gamma_1| \to |\Gamma_2|$.

Now we clarify the meaning of some standard terms which may have ambiguous interpretations in this topological context. Let $G = (\Gamma, \varphi)$ be a Σ-graph and let $e \in E(\Gamma)$. By Γ/e we mean the graph obtained by identifying the end vertices of e and deleting (only) the edge e. Thus edge contractions can create parallel edges and/or loops. By G/e we mean the surface

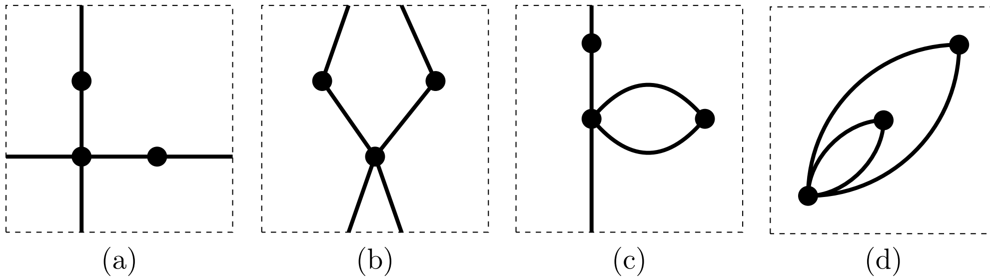

FIGURE 1. Four pairwise non isomorphic torus graphs that all have the same underlying graph. Here and in subsequent diagrams we use the standard representation of the torus as a square with opposite edges identified appropriately. In (a) there is just one face which is a degenerate cellular face of degree 8. In (b) there are two faces. One is a degenerate quadrilateral (i.e degree 4) cellular face. The other is a degenerate face of genus 0 with two boundary walks, both of length 2. In (c) there is a non degenerate digon (i.e cellular of degree 2) face and a non cellular genus zero degenerate face. In this case the non cellular face has two boundary walks of length 2 and 4 respectively. In (d) there are three faces: a non degenerate digon, a degenerate quadrilateral face and a face of genus 1. We also observe that (d) is an inessential torus graph, while both (b) and (c) are annular and (a) is essential but not annular (see Section 6 for the relevant definitions).

graph obtained by collapsing the arc corresponding to e to a single point. Clearly the underlying graph of G/e is Γ/e. A face, F, of G is a connected component of $\Sigma - \varphi(|\Gamma|)$. A cellular face is one that is homeomorphic to $\mathbb{R}^2$. A surface graph is cellular if all its faces are cellular.

Associated to any face F there is a collection of closed walks, one for each topological end of F, called the boundary walks of F. We outline the construction of these walks, for full details see Chapter 4 of [15]. Our setting is slightly different in that we allow our faces to be non cellular, however the discussion in loc. cit. can be adapted to our setting with the understanding that a face may now have more than one end and thus more than one boundary walk.

It is well known that associated to a surface graph (Γ, φ) there is a rotation system on the graph Γ. Given an end E of F, we choose a half edge d that lies in the boundary of E. Now the rotation system corresponding to G associates two other half edges d', d'' to d as follows: d, d' are half edges in the same edge and d, d'' are half edges that have a common vertex and both lie in the closure of the end E. The sequence d', d, d'' is part of the boundary walk associated to E. By iterating this construction on the sequence of half edges, we ultimately obtain a closed boundary walk associated to the end E. We note that the boundary walks are only unique up to cyclic permutation and choice of orientation.

We say that F is non degenerate if no vertex occurs more than once in the set of boundary walks of F. For a cellular face F, the degree of F, denoted $|F|$, is the edge length of its unique boundary walk, which of course may differ from the number of vertices or edges embedded in the boundary of the face in degenerate cases. We write f_i for the number of cellular faces of degree i. Note that if Σ is connected then $f_0 = 1$ if Σ is a sphere and Γ comprises a single vertex, and $f_0 = 0$ otherwise. See Figure 1 for some examples of torus graphs illustrating some of the definitions mentioned in this section.

By the Heffter-Edmonds-Ringel rotation principle [15, pp. 90–91], the surface Σ and the Σ-graph (Γ, φ) are determined up to isomorphism by data consisting of a rotation system on Γ, a partition of the set of boundary walks associated to the rotation system and for each part of that partition a nonnegative integer that represents the genus of the corresponding facial region. Note that we do not assume that our surface graphs are cellular, hence the necessity for the partition of the set of facial walks and data on the facial genus. See [15] for details of this. We will be interested in determining certain classes of surface graphs up to isomorphism. The above-mentioned principle allows us to argue topologically using properties of surfaces and curves in surfaces to deduce combinatorial information and we choose to write our arguments using topological terminology based on this.

Finally we note that we extend much of the standard language of graph theory concerning subgraphs, intersections and unions to surface graphs, understanding that these terms apply to the underlying graphs. Thus if $G = (\Gamma, \varphi)$ is a Σ-graph, a Σ-subgraph of G is a pair $(\Gamma', \varphi|_{\Gamma'})$ where Γ' is a subgraph of Γ and $\varphi|_{\Gamma'}$ is the restriction of φ to $|\Gamma'|$. If H_1, H_2 are Σ-subgraphs of G then $H_1 \cup H_2$, respectively $H_1 \cap H_2$, is the Σ-graph whose underlying graph is the union, respectively intersection, of the underlying graphs of H_1 and H_2.

3. Sparsity

For a graph $\Gamma = (V, E)$ define $\gamma(\Gamma) = 2|V| - |E|$. For $l \leq 2$ we say that Γ is $(2, l)$-sparse (or just sparse if l is clear from the context) if, $\gamma(\Gamma') \geq l$ for every nonempty subgraph Γ' of Γ. We say that Γ is $(2, l)$-tight if it is $(2, l)$-sparse and $\gamma(\Gamma) = l$. We will be particularly interested in $(2, 2)$-sparse graphs. Note that $(2, 2)$-tight graphs cannot have loop edges but can have parallel edges.

We record some standard elementary facts for later use. The proofs are straightforward and we omit them. Suppose that Γ_1, Γ_2 are subgraphs of Γ. Then

$$\gamma(\Gamma_1 \cup \Gamma_2) = \gamma(\Gamma_1) + \gamma(\Gamma_2) - \gamma(\Gamma_1 \cap \Gamma_2) \tag{1}$$

Lemma 3.1. *Suppose that Γ is $(2, 2)$-sparse and that $\gamma(\Gamma') \leq 3$ for some subgraph Γ' of Γ. Then Γ' is connected.* □

Lemma 3.2. *Suppose that Γ_1, Γ_2 are $(2, 2)$-tight subgraphs of a $(2, 2)$-sparse graph Γ. If $\Gamma_1 \cap \Gamma_2$ is not empty then both $\Gamma_1 \cup \Gamma_2$ and $\Gamma_1 \cap \Gamma_2$ are $(2, 2)$-tight.* □

We also record a straightforward consequence of Euler's polyhedral formula.

Theorem 3.3. *If Σ is a connected boundaryless compact orientable surface of genus g and G is a cellular Σ-graph then*

$$\sum_{i \geq 0} (4 - i) f_i = 8 - 8g - 2\gamma(G) \tag{2}$$

Proof. Use the polyhedral formula and the fact that $\sum i f_i = 2|E|$. □

Next we will derive an elementary but subsequently very useful principle relating the function γ, which is defined purely in terms of the underlying graph, to the embedding of the graph in the surface. We note that related results and special cases of this have appeared elsewhere (notably [3, 4, 7]). Here we attempt to express this principle in more general terms: for surfaces of arbitrary genus and for a variety of sparsity counts.

In order to state the result we introduce some notation. Let H be a Σ-subgraph of the Σ-graph G and suppose that F is a face of H. Let $\mathrm{int}_G(F)$ be the Σ-subgraph of G consisting of all vertices and edges of G that lie inside $\overline{F}$ (the topological closure of F in Σ). Let $\mathrm{ext}_G(F)$ be the Σ-subgraph of G consisting of all vertices and edge of G that lie in $\Sigma - F$. Define $\partial F = \mathrm{int}_G(F) \cap \mathrm{ext}_G(F)$ and note that ∂F is the smallest Σ-subgraph of G that supports the boundary walks of F. It follows that $G = \mathrm{int}_G(F) \cup \mathrm{ext}_G(F)$ since any edge joining $\mathrm{int}_G(F)$ to $\mathrm{ext}_G(F)$ must pass through ∂F.

Theorem 3.4. *Suppose that $l \leq 2$ and that G is a $(2,l)$-tight Σ-graph. If H is a Σ-subgraph of G and F is a face of H, then $\gamma(H \cup \mathrm{int}_G(F)) \leq \gamma(H)$.*

Proof. By (1) we have

$$\gamma(H \cup \mathrm{int}_G(F)) = \gamma(H) + \gamma(\mathrm{int}_G(F)) - \gamma(H \cap \mathrm{int}_G(F)).$$

Now, $H \cap \mathrm{int}_G(F) = \mathrm{ext}_G(F) \cap \mathrm{int}_G(F)$ and using (1) again, we see that

$$\begin{aligned}
\gamma(H \cap \mathrm{int}_G(F)) &= \gamma(\mathrm{ext}_G(F) \cap \mathrm{int}_G(F)) \\
&= \gamma(\mathrm{int}_G(F)) + \gamma(\mathrm{ext}_G(F)) - \gamma(\mathrm{ext}_G(F) \cup \mathrm{int}_G(F)) \\
&= \gamma(\mathrm{int}_G(F)) + \gamma(\mathrm{ext}_G(F)) - \gamma(G) \\
&= \gamma(\mathrm{int}_G(F)) + \gamma(\mathrm{ext}_G(F)) - l \\
&\geq \gamma(\mathrm{int}_G(F)).
\end{aligned}$$

The last inequality above follows from applying the sparsity of G to the nonempty Σ-subgraph $\mathrm{ext}_G(F)$. ☐

Corollary 3.5. *Suppose that $l \leq 2$ and that G is a $(2,l)$-tight Σ-graph. If H is a Σ-subgraph of G and F is a face of H, then $\gamma(\mathrm{ext}_G(F)) \leq \gamma(H)$.*

Proof. Let $J_1, \cdots, J_k$ be all the faces of H that are different from F. Then $\mathrm{ext}_G(F) = H \cup \bigcup_{i=1}^k \mathrm{int}_G(J_i)$. Now the conclusion follows from repeated applications of Theorem 3.4. ☐

We conclude this section by making some straightforward observations about the Σ-subgraphs $\mathrm{int}_G(F)$ and $\mathrm{ext}_G(F)$ that will be useful in the sequel. Any face of $\mathrm{int}_G(F)$ that is contained in F is also a face of G. On the other hand, there are one or more faces of $\mathrm{int}_G(F)$ which are contained in $\Sigma - \overline{F}$. We call such a face an external face of $\mathrm{int}_G(F)$. Such an external face need not be a face of G. Note that if F has a unique boundary walk that is a simple cycle, then $\mathrm{int}_G(F)$ has just one external face. In general it may have more than one external face. Observe that $\mathrm{ext}_G(F)$ has one exceptional face, namely F, such that all other faces of $\mathrm{ext}_G(F)$ are also faces of G.

4. Inductive operations on surface graphs

In this section we will focus on topological inductive operations on graphs that are natural in the context of $(2,l)$-tight graphs. We consider three types of contractions associated to cellular faces of degree 2, 3 and 4, respectively called digons, triangles and quadrilaterals hereafter. In each case the contraction decreases the number of vertices by one and the number of edges by two. We investigate necessary and sufficient conditions for these moves to preserve the property of being $(2,2)$-tight. We pay particular attention to degenerate cases as these play an important role later.

4.1. **Digon and triangle contractions.** Let G be a Σ-graph and suppose that D is a digon of G with boundary walk v_1, e_1, v_2, e_2, v_1 such that $v_1 \neq v_2$ and $e_1 \neq e_2$. Let $G_D = (G/e_1) - e_2$. Observe that $(G/e_1) - e_2$ is canonically isomorphic to $(G/e_2) - e_1$, so G_D depends only on the digon and not the particular choice of labelling of the edges. We remark that, for a connected surface Σ, while a digon in a $(2,2)$-sparse Σ-graph necessarily has distinct vertices, it may have a degenerate boundary, but only in the case that the graph is a single (non loop) edge and Σ is a sphere.

The proof of the following is straightforward and we omit it.

Lemma 4.1. *G is $(2,l)$-sparse if and only if G_D is $(2,l)$-sparse*

Now suppose that T is a triangle in G with boundary walk $v_1, e_1, v_2, e_2, v_3, e_3, v_1$ such that $v_1 \neq v_2$ and $e_1 \neq e_2$. Let $G_{T,e_1} = (G/e_1) - e_2$. Again we omit the proof of the following lemma as it is a straightforward consequence of the definitions.

Lemma 4.2. *Suppose that G is $(2,l)$-sparse and that G_{T,e_1} is not $(2,l)$-sparse. Then there is a Σ-subgraph H of G that contains e_1 but not v_3 such that $\gamma(H) = l$*

We refer to the graph H whose existence is asserted in Lemma 4.2 as a blocker for the contraction G_{T,e_1}.

We note that a triangle in a $(2,2)$-sparse surface graph necessarily has a non degenerate boundary walk since any degeneracy would entail a (forbidden) loop edge. Thus, in this case there are three possible contractions (one for each of the edges) associated to any such face.

Lemma 4.3. *Suppose that G is a $(2,2)$-sparse Σ-graph and that T is a triangle with edges e_1, e_2, e_3. Then at least two of the Σ-graphs $G_{T,e_1}, G_{T,e_2}, G_{T,e_3}$ are $(2,2)$-sparse.*

Proof. Suppose that there are blockers H_1, respectively H_2, for G_{T,e_1} respectively G_{T,e_2}. Then $v_1, v_3 \in H_1 \cup H_2$. However $v_3 \notin H_1$ and $v_1 \notin H_2$ so $e_3 \notin H_1 \cup H_2$. However $v_2 \in H_1 \cap H_2$ so by Lemma 3.2, $H_1 \cup H_2$ is $(2,2)$-tight. This contradicts the sparsity of G. $\square$

4.2. **Quadrilateral contractions.** In the case of quadrilaterals we consider a somewhat different contraction move. In this case the analysis is a little more complicated and we include the details.

Suppose that Q is a quadrilateral of G with possibly degenerate boundary walk $v_1, e_1, v_2, e_2, v_3, e_3, v_4, e_4, v_1$. Suppose that $v_1 \neq v_3$ and $e_1 \neq e_3$. Let d be a new edge that joins v_1 and v_3 and is embedded as a diagonal of the quadrilateral Q. Define G_{Q,v_1,v_3} to be $(G \cup \{d\})/d - \{e_1, e_3\}$. Clearly the underlying graph of G_{Q,v_1,v_3} is obtained from Γ by identifying the vertices v_1 and v_3 and then deleting e_1 and e_3. Thus $\gamma(G) = \gamma(G_{Q,v_1,v_3})$. However this quadrilateral contraction move does not necessarily preserve $(2,l)$-sparsity.

Lemma 4.4. *Suppose that G is $(2,l)$-sparse but G_{Q,v_1,v_3} is not $(2,l)$-sparse. Then at least one of the following statements is true.*

- *(i) There is some Σ-subgraph H of G such that $v_1, v_3 \in H$, exactly one of v_2, v_4 is in H and $\gamma(H) = l$. (H is called a type 1 blocker.)*
- *(ii) There is some Σ-subgraph K of G such that $v_1, v_3 \in K$, $v_2, v_4 \notin K$ and $\gamma(K) = l + 1$. (K is called a type 2 blocker.)*

Proof. Let K be a maximal Σ-subgraph of G_{Q,v_1,v_3} satisfying $\gamma(K) \leq l - 1$. Let z be the vertex of G_{Q,v_1,v_3} corresponding to v_1 and v_3. Clearly $z \in K$, otherwise K would also be a Σ-subgraph of G. Let H be the maximal Σ-subgraph of G satisfying $(H \cup \{d\})/d - \{e_1, e_3\} = K$.

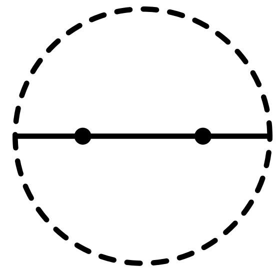

FIGURE 2. A $(2,2)$-tight projective plane graph. Here we are using the representation of the projective plane as a disc with antipodal boundary points identified. This surface graph has a single quadrilateral face, with a degenerate boundary walk.

It is clear that H is an induced Σ-subgraph of G, since K is an induced Σ-subgraph. If $\{v_2, v_4\} \subset H$, then $\gamma(H) = \gamma(K) \leq l - 1$ which contradicts the sparsity of G. So at most one of v_2, v_4 belongs to H. Also, it is clear that $l \leq \gamma(H) \leq \gamma(K) + 2 \leq l + 1$. So $\gamma(H) = l$ or $l + 1$. If $\gamma(H) = l$ and one of $v_2, v_4 \in H$ then (i) is true. If $\gamma(H) = l$ and neither of v_2, v_4 is in H, then let $H' = H \cup \{v_2\} \cup \{e_1, e_2\}$. Now observe that $e_1 \neq e_2$ since $v_1 \neq v_3$. Thus $\gamma(H') = \gamma(H) = l$ and, again, (i) is true. Finally if $\gamma(H) = l + 1$. Then $\gamma(H) = \gamma(K) + 2$ and since H is an induced graph, it follows that neither of v_2, v_4 belongs to H. Thus (ii) is true in this case. □

In the special case that $l = 2$, certain degenerate cases cannot arise.

Lemma 4.5. *Suppose that G is a $(2,2)$-sparse Σ-graph and that Q is a quadrilateral face of G with boundary walk $v_1, e_1, v_2, e_2, v_3, e_3, v_4, e_4, v_1$. For $i = 1, 2, 3$, $v_i \neq v_{i+1}$, and $v_1 \neq v_4$.*

Proof. Loop edges are forbidden in a $(2,2)$-sparse graph. □

Lemma 4.6. *Suppose that Σ is orientable, G is a $(2,2)$-tight Σ-graph and that Q is a quadrilateral face of G with boundary walk $v_1, e_1, v_2, e_2, v_3, e_3, v_4, e_4, v_1$.*

- *(i) $e_i \neq e_j$ for $1 \leq i < j \leq 4$.*
- *(ii) If $v_1 = v_3$ then $v_2 \neq v_4$ and G_{Q,v_2,v_4} is $(2,2)$-tight.*

Proof. We observe that since Σ is orientable, a repeated edge in ∂Q implies the existence either of a vertex of degree one or of a loop edge. Both of these are forbidden in a $(2,2)$-tight graph. This proves (i).

Now suppose that $v_1 = v_3$. If $v_2 = v_4$ then by Lemma 4.5 and the sparsity of G, ∂Q has exactly two vertices and two edges. This contradicts (ii), thus $v_2 \neq v_4$. Suppose that G_{Q,v_2,v_4} is not $(2,2)$-tight. By Lemma 4.4 there is a blocker for this contraction. Since $v_1 = v_3$ by assumption, the blocker must be a type 2 blocker. Thus we have a Σ-subgraph K such that $\gamma(K) = 3$, $v_2, v_4 \in K$ and $v_1 \notin K$. However, by Lemma 4.6 there are at least four edges joining v_1 to K, contradicting the sparsity of G and completing the proof of (ii). □

See Figure 2 for an example of $(2,2)$-tight projective plane graph whose only face is a quadrilateral with repeated edges in the boundary walk. This example shows that orientability is a necessary hypothesis in the statement of Lemma 4.6.

Lemma 4.7. *Suppose that Σ is orientable, G is $(2,2)$-tight and Q is a quadrilateral face of G such that neither G_{Q,v_1,v_3} nor G_{Q,v_2,v_4} is $(2,2)$-sparse. Then Q has a non degenerate boundary.*

Furthermore, if H_1 and H_2 are blockers for G_{Q,v_1,v_3} respectively G_{Q,v_2,v_4}, then both H_1 and H_2 are type 2 blockers and $H_1 \cap H_2 = \emptyset$.

Proof. The non degeneracy of the boundary walk of Q follows immediately from Lemmas 4.5 and 4.6.

Now suppose that one of the blockers, say H_1, is of type 1 and suppose that $v_2 \notin H_1$. Then $v_4 \in H_1 \cap H_2$. So $\gamma(H_1 \cup H_2) = \gamma(H_1) + \gamma(H_2) - \gamma(H_1 \cap H_2) \leq 2 + \gamma(H_2) - 2 = \gamma(H_2)$. Now if H_2 is also type 1 then $\gamma(H_1 \cup H_2) = 2$. However $v_1, v_2, v_3 \in H_1 \cup H_2$ but $H_1 \cup H_2$ does not contain one of e_1, e_2 which contradicts the sparsity of G. Similarly if H_2 is type 2, then $\gamma(H_1 \cup H_2) \leq 3$, but $H_1 \cup H_2$ does not contain either of e_1, e_2, again contradicting the sparsity of G.

So both H_1 and H_2 are type 2 blockers. Moreover $v_1, v_2, v_3, v_4 \in H_1 \cup H_2$ but $e_1, e_2, e_3, e_4 \notin H_1 \cup H_2$. Now

$$\begin{aligned} 2 &\leq \gamma(H_1 \cup H_2 \cup \{e_1, e_2, e_3, e_4\}) \\ &= \gamma(H_1) + \gamma(H_2) - \gamma(H_1 \cap H_2) - 4 \\ &= 2 - \gamma(H_1 \cap H_2) \end{aligned}$$

So $\gamma(H_1 \cap H_2) \leq 0$ which implies that $H_1 \cap H_2 = \emptyset$. □

4.3. **Simple loops in surfaces.** Now we briefly digress to review some necessary terminology and facts from low dimensional topology. Proofs of all of the assertions below can be found in (or at least easily deduced from) many sources (for example [6]). A loop in a surface Σ is a continuous function $\alpha : S^1 \to \Sigma$. We say that α is simple if it is injective. We say that α is non separating if $\Sigma - \alpha(S^1)$ has the same number of connected components as Σ. Given simple loops α, β in Σ, recall that the geometric intersection number is defined by

$$i(\alpha, \beta) = \min |\alpha'(S^1) \cap \beta'(S^1)|$$

where α', respectively β', varies over all simple loops that are homotopic to α, respectively β. If $i(\alpha, \beta) \neq 0$ then both α and β are essential: that is to say they are not null homotopic. If $i(\alpha, \beta) = 1$ then both α and β are non separating in Σ. In the special case that Σ is the torus, if $i(\alpha, \beta) = 0$ and $i(\beta, \delta) = 0$ then $i(\alpha, \delta) = 0$.

Given a simple loop α in a surface Σ we say that $\Sigma - \alpha(S^1)$ is the surface obtained by cutting along α. Given a surface Σ with boundary we can cap a boundary component by gluing a copy of a closed disc to the surface along the given boundary component.

If Σ is an orientable surface of genus g and α is a non separating simple loop in Σ then we form Σ^α by removing a tubular neighbourhood of α [1] and then capping the two resulting new boundary components by gluing a copy of the disc into each of the resulting open ends. Clearly Σ^α is an orientable surface of genus $g - 1$. We will refer to this process as cutting and capping along α.

Suppose that G is a Σ-graph and let F be a face of G. Further suppose that α is a non separating loop in Σ such that $\alpha(S^1) \subset F$. By cutting and capping Σ along α we can form a Σ^α-graph, denoted G^α, which has the same underlying graph as G. Observe that all faces of G^α except the one(s) corresponding to F are also faces of G.

Finally some terminology. If $G = (\Gamma, \varphi)$ is a Σ-graph and α is a loop in Σ, we say that α is contained in G if $\alpha(S^1) \subset \varphi(|\Gamma|)$.

[1] that is a neighbourhood of α that is homeomorphic to an open annulus

Now we return to the situation of Lemma 4.7. Suppose that Q is a quadrilateral in G as in the statement of that lemma. We say that the quadrilateral Q is blocked. By Lemma 3.1 the blocker H_1 is connected so it is possible to find a simple walk from v_1 to v_3 in H_1. By concatenating the geometric realisation of this walk with the diagonal of Q joining v_3 and v_1 we obtain a simple loop in Σ, which we denote by α_1. Note that we can choose different parameterisations of this loop, but this ambiguity will make no difference in our context. Similarly we construct another simple loop, denoted α_2, by concatenating a walk in H_2 with the diagonal of Q that joins v_4 and v_2. Now since $H_1 \cap H_2$ is empty by Lemma 4.7, we can choose these loops so that they intersect transversely at exactly one point (where the diagonals meet). Thus these loops have geometric intersection number equal to one. In particular, we note that both α_1 and α_2 must be non separating loops in Σ. These loops will play an important role in the following sections.

5. Irreducible surface graphs

Let G be a $(2,2)$-tight Σ-graph. In light of Lemmas 4.1 and 4.3 we say that G is irreducible if it has no digons, no triangles and if, for every quadrilateral face of G, both of the possible contractions result in graphs that are not $(2,2)$-sparse.

For each of the contractions described in Section 4 there are the corresponding vertex splitting moves. More precisely, if $G' = G_D$, respectively $G' = G_{T,e}$, respectively $G' = G_{Q,u,v}$ for some digon D, respectively triangle T and edge $e \in \partial T$, respectively quadrilateral Q and vertices $u, v \in \partial Q$, then we say that G is obtained from G' by a digon, respectively triangle, respectively quadrilateral split. Thus every $(2,2)$-tight Σ-graph can be constructed from some irreducible by applying a sequence of digon/triangle/quadrilateral splits. Our goal is to identify, for various surfaces, the set of irreducibles.

Conjecture 5.1. *If Σ is a surface with finite genus and finitely many boundary components and punctures, then there are finitely many distinct isomorphism classes of irreducible $(2,2)$-tight Σ-graphs.*

We will address some special cases of Conjecture 5.1 in this and later sections. Let $\mathbb{S}$ be the 2-sphere.

Theorem 5.2. *If G is a $(2,2)$-tight $\mathbb{S}$-graph with at least two vertices then G has at least two faces of degree at most 3. In particular, any $(2,2)$-tight $\mathbb{S}$-graph can be constructed from a single vertex by a sequence of digon and/or triangle splits.*

Proof. By Lemma 3.1, G is connected and therefore cellular. Since G has at least two vertices, $f_0 = 0$. Also $f_1 = 0$ by sparsity, so by Theorem 3.3, we see that $2f_2 + f_3 \geq 4$. □

The case of plane graphs is similarly straightforward.

Corollary 5.3. *If G is a $(2,2)$-tight $\mathbb{R}^2$-graph with at least two vertices then G has at least one cellular face of degree at most 3. In particular, any $(2,2)$-tight $\mathbb{R}^2$-graph can be constructed from a single vertex by a sequence of digon and/or triangle splits.*

Proof. Cap (i.e fill in the puncture of) the non cellular face of G and then apply Theorem 5.2. □

Theorem 5.2 and Corollary 5.3 are implicit already in other places in the literature, we include them here for completeness. We note that in both cases the quadrilateral splitting move is not required in the inductive characterisation.

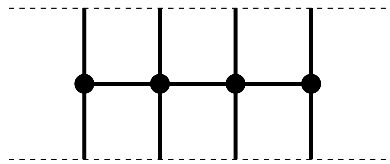

Figure 3. The top and bottom dashed lines are identified to form an open cylinder. Thus this diagram represents an $\mathbb{A}$-graph with no digons or triangular faces and 3 quadrilateral faces.

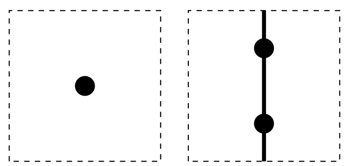

Figure 4. The two non cellular irreducible torus graphs. Note that by cutting the torus along a non separating loop these graphs can also be viewed as graphs in the twice punctured sphere.

Now let $\mathbb{A}$ be the twice punctured sphere $\mathbb{R}^2 - \{(0,0)\}$. Observe that for any positive integer n, it is straightforward to construct an $\mathbb{A}$-graph that has no digons or triangles, but has n quadrilateral faces (see Figure 2). So, in contrast to the cases of the sphere or plane, we do require the quadrilateral contraction move in order to have finitely many irreducible $(2,2)$-tight $\mathbb{A}$-graphs.

There are two obvious examples of irreducible $(2,2)$-tight $\mathbb{A}$-graphs, with one vertex and two vertices respectively: see Figure 4.

Theorem 5.4. *If G is an irreducible $(2,2)$-tight $\mathbb{A}$-graph, then G is isomorphic to one of the $\mathbb{A}$-graphs shown in Figure 4.*

Proof. There are two cases to consider. First suppose that G does not separate the two punctures of $\mathbb{A}$. In this case it follows easily from Theorem 5.2 that if G has at least two vertices then it has at least one triangular face and so is not irreducible.

Now suppose that G does separate the punctures of $\mathbb{A}$. Clearly G has exactly two non cellular faces. By capping these two faces, we create a $(2,2)$-tight $\mathbb{S}$-graph $\tilde{G}$. This graph satisfies $2f_2 + f_3 = 4 + f_5 + 2f_6 + \cdots$ and since all but two of the faces of $\tilde{G}$ are also faces of the irreducible G, it follows that the two exceptional faces of $\tilde{G}$ are digons and all other faces are quadrilateral faces of G. Thus it suffices to show that there cannot be any quadrilateral faces in G.

For a contradiction, suppose that Q is a quadrilateral. Since G is irreducible, both possible contractions of Q are blocked and we infer the existence of simple loops α_1 and α_2 as described at the end of Section 4. Recall that these loops intersect transversely at exactly one point and thus α_1 is non separating in $\mathbb{A}$. However the Jordan Curve Theorem tells us that any simple loop in $\mathbb{A}$ must be separating. □

6. Subgraphs of irreducibles

Recall that a surface is closed if it is compact and without boundary. The goal of this section is to show that if Σ is an orientable closed surface and $G = (\Gamma, \varphi)$ is an irreducible $(2,2)$-tight Σ-graph, then any $(2,2)$-tight Σ-subgraph of G is also irreducible.

Let $H = (\Lambda, \varphi|_{|\Lambda|})$ be a Σ-subgraph of G. We say that H is inessential if there is some embedded open disc $U \subset \Sigma$ such that $\varphi(|\Lambda|) \subset U$. If there is no such disc then H is essential. Observe that if F is a cellular face of G that has a non degenerate boundary walk, then ∂F is inessential: let U be an open disc neighbourhood of the embedded closed disc $\overline{F}$. We also note that if H is inessential and connected then it has at most one non cellular face F. Moreover if we cut and cap along an appropriate[2] set of loops in F we obtain an $\mathbb{S}$-graph which, in this section, we will denote by $\hat{H}$.

Let K_1 be the graph with one vertex and no edges. Let K_2 be the complete graph on two vertices. For $n \geq 2$ let C_n be the n-cycle graph (in particular C_2 has exactly two parallel edges).

Lemma 6.1. *Suppose that Σ is an orientable closed surface, $G = (\Gamma, \varphi)$ is an irreducible $(2,2)$-tight Σ-graph and H is a Σ-subgraph of G whose underlying graph is isomorphic to either C_2 or C_3. Then H is essential.*

Proof. Suppose that the underlying graph of H is isomorphic to C_2. The other case is similar. Suppose that H is inessential. Let U be an open disc that contains $\varphi(\Lambda)$. Clearly there is a digon face D of H that is contained in U. Now let K be the $\mathbb{S}$-graph obtained by cutting and capping the external face of $\mathrm{int}_G(D)$. By Theorem 3.4, $\gamma(K) = 2$ and by Theorem 3.3, K has at least two faces of degree at most 3. One of these faces is also a face of G contradicting the irreducibility of G. $\square$

Lemma 6.2. *Suppose that Σ is an orientable closed surface, $G = (\Gamma, \varphi)$ is an irreducible $(2,2)$-tight Σ-graph and H is an inessential Σ-subgraph of G and that $\gamma(H) = 2$. Then the underlying graph of H is K_1.*

Proof. Suppose that H has at least two vertices. Then by Theorem 3.3, $\hat{H}$ has at least two faces of degree at most 3. If one of these is a triangle or a digon with non degenerate boundary then the underlying graph of H contains a copy of C_2 or C_3 which contradicts Lemma 6.1. Therefore $\hat{H}$ must have two digon faces both of which have degenerate boundaries. However, as pointed out in Section 4.1, no $\mathbb{S}$-graph can have more than one degenerate digon. $\square$

Lemma 6.3. *Suppose that Σ is an orientable closed surface, $G = (\Gamma, \varphi)$ is an irreducible $(2,2)$-tight Σ-graph, H is an inessential Σ-subgraph of G and that $\gamma(H) = 3$. Then the underlying graph of H is K_2.*

Proof. By Theorem 3.3, $\hat{H}$ satisfies $2f_2 + f_3 = 2 + f_5 + 2f_6 + \cdots$. As in the proof of Lemma 6.2 we see that $\hat{H}$ cannot have a triangle or a digon with non degenerate boundary. So the only possibility is that $\hat{H}$ has a digon face with degenerate boundary. As pointed out in Section 4.1, there is only one $\mathbb{S}$-graph with a degenerate digon face and its underlying graph is indeed K_2. $\square$

The case of a Σ-subgraph isomorphic to C_4 is a little more involved.

[2] more precisely, a maximal set of pairwise disjoint simple loops that is non separating in F

Lemma 6.4. *Suppose that Σ is an orientable closed surface, $G = (\Gamma, \varphi)$ is an irreducible $(2,2)$-tight Σ-graph and H is an inessential subgraph of G whose underlying graph is isomorphic to C_4. Then H is the boundary of some quadrilateral face of G.*

Proof. Suppose that U is an embedded disc containing $\varphi(|\Lambda|)$ and let R be the face of H that is contained in U. First observe that $\gamma(H) = 4$, so by Theorem 3.4, $\gamma(\mathrm{int}_G(R)) \leq 4$. Now, by Lemma 6.1, $\mathrm{int}_G(R)$ has no digons or triangles and it follows easily from Theorem 3.3 that $\gamma(\mathrm{int}_G(R)) = 4$ and that all the cellular faces of $\mathrm{int}_G(R)$ are quadrilaterals: that is to say that $\mathrm{int}_G(R)$ is in fact a quadrangulation of $\overline{R}$.

Now, let Q (with boundary vertices v_1, v_2, v_3, v_4) be a quadrilateral face of $\mathrm{int}_G(R)$ that is contained in R. Since G is irreducible, we have blockers H_1 and H_2 for the two possible contractions of Q, as described in Lemma 4.7. Also we have simple loops α_1 and α_2 as described in Section 4. These loops intersect transversely at one point in Q. If w_1, w_2, w_3, w_4 are the vertices of ∂R in cyclic order, it follows that one of the loops, say α_1, contains w_1 and w_3 and that α_2 contains w_2 and w_4. Thus α_2 divides R into disjoint open subsets R_1 and R_3 (see Figure 5) where $w_1, v_1 \in \overline{R_1}$ and $w_3, v_3 \in \overline{R_3}$. Now we can decompose the blocker H_1 as $K_e \cup K_1 \cup K_3$, where $K_e = \mathrm{ext}_G(R) \cap H_1$, K_1 is the part of H_1 contained in $\overline{R}_1$ and K_3 is the part of H_1 contained in $\overline{R}_3$. It is clear that $K_e \cap K_1 = \{w_1\}$ and $K_e \cap K_3 = \{w_3\}$. Therefore, by (1),

$$3 = \gamma(H_1) = \gamma(K_e) + \gamma(K_1) + \gamma(K_3) - 4.$$

Using the sparsity of G it follows that at least one of $\gamma(K_1)$ or $\gamma(K_3)$ is equal to 2. Now K_1 and K_3 are both inessential Σ-subgraphs of G since $\overline{R}_1$ and $\overline{R}_3$ are both embedded closed discs in Σ. It follows from Lemma 6.2 that at least one of K_1 or K_3 is a single vertex. So either $v_1 = w_1$ or $v_3 = w_3$. We have shown that at least one of v_1 or v_3 actually lies in the boundary of R. Similarly at least one of v_2 or v_4 lies in the boundary of R.

Thus we have shown that if Q is any quadrilateral face of G contained in R then ∂Q and ∂R share at least one edge. Now it is an elementary exercise to show that in any quadrangulation of R that has this property, either there are no quadrilaterals properly contained in R, or some quadrilateral has a boundary vertex with degree 2. Clearly, by Lemma 4.7, no quadrilateral face of the irreducible graph G can have a boundary vertex of degree 2. It follows that there are no quadrilateral faces of G that are properly contained in R and so R is itself a face of G. □

We say that a Σ-subgraph $H = (\Lambda, \varphi|_{|\Lambda|})$ of G is annular if it is essential and $\varphi(|\Lambda|)$ is contained in some embedded open annulus of Σ. [3]

Let $\mathfrak{B}$ be the (unique) $(2,2)$-tight graph with 3 vertices, one of which has degree 4. See Figure 1 for some examples of annular and non annular embeddings of $\mathfrak{B}$ in the torus.

Lemma 6.5. *Suppose that Σ is an orientable closed surface, $G = (\Gamma, \varphi)$ is an irreducible $(2,2)$-tight Σ-graph and $H = (\Lambda, \varphi|_{|\Lambda|})$ is a Σ-subgraph of G whose underlying graph is isomorphic to $\mathfrak{B}$. Then H is not annular.*

Proof. Suppose, seeking a contradiction, that H is annular. Let g be the genus of Σ. Since H is annular, $g \geq 1$. Now we observe that H must be obtained from one of the torus graphs in Figure 1 (b) or (c) by adding $g-1$ handles to the non cellular face and possibly cutting and capping the resulting face along a separating curve. In case (c) H has a non degenerate

[3] We note that this is equivalent to saying that image of the induced homomorphism of fundamental groups $\varphi_* : \pi_1(|\Lambda|) \to \pi_1(\Sigma)$ is nontrivial and cyclic. We will not need this fact, so we omit the proof.

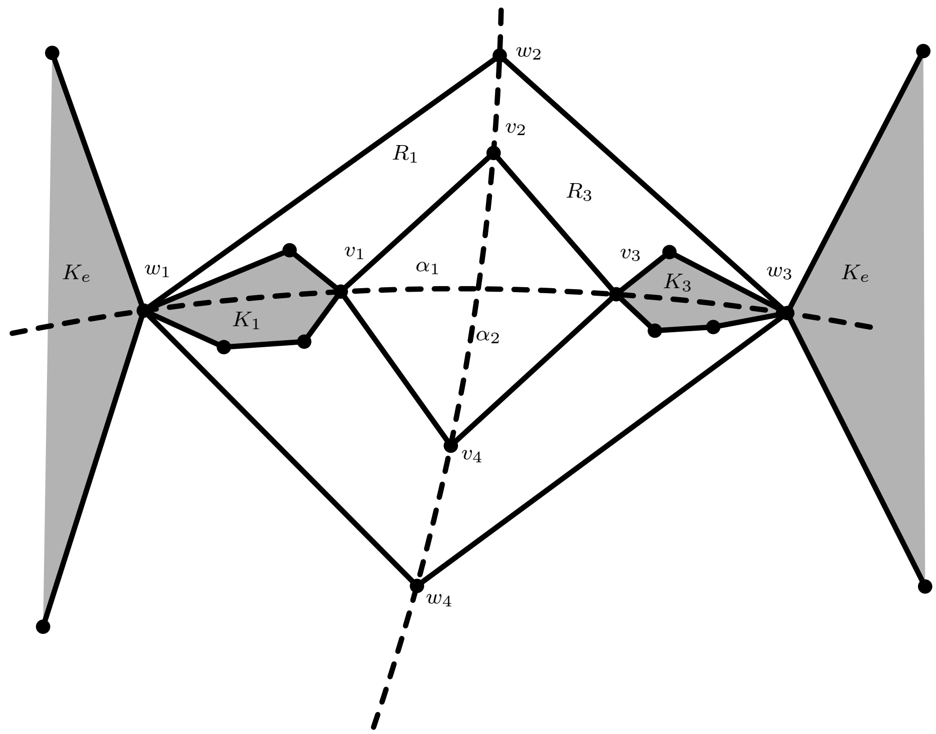

FIGURE 5. From the proof of Lemma 6.4: the shaded region represents the blocker for the contraction G_{Q,v_1,v_3}.

digon face, contradicting Lemma 6.1. Thus H is isomorphic to the surface graph obtained from Figure 1 (b) as described above. Let U be an open annulus containing $\varphi(|\Lambda|)$ and let R be the (unique) face of H that is contained in U. Observe that $\gamma(H) = 2$, so by Theorem 3.4, $\gamma(\text{int}_G(R)) = 2$. Let K be the $\mathbb{S}$-graph obtained by cutting and capping the external faces of $\text{int}_G(R)$ (there could be more than one in this case). Now K is a $(2,2)$-tight $\mathbb{S}$-graph with two digon faces. Since all other faces of K are also faces of the irreducible G, it follows easily from Theorem 3.3 that all other faces of K are quadrilaterals. Thus, all faces of G that are contained in R are in fact quadrilaterals.

Now we can argue, using a straightforward modification of the argument from the proof of Lemma 6.4, that any quadrilateral face of G that is contained in R must in fact share a boundary edge with R. Again, following the proof of Lemma 6.4 it follows that R itself must be a face of G. However this contradicts Lemma 4.7 where we showed that any quadrilateral face of an irreducible has a non degenerate boundary. □

Now the main result of this section: a tight Σ-subgraph of an irreducible is also irreducible.

Theorem 6.6. *Suppose that Σ is an orientable closed surface, $G = (\Gamma, \varphi)$ is an irreducible $(2,2)$-tight Σ-graph and Λ is a $(2,2)$-tight Σ-subgraph of Γ. Then $H = (\Lambda, \varphi|_{|\Lambda|})$ is an irreducible Σ-graph.*

Proof. We see that H cannot have any triangle or digon: the boundary of such a face must be non degenerate by $(2,2)$-sparsity, and so would be an inessential subgraph of G that would contradict Lemma 6.1. Now suppose that Q is a quadrilateral face of H. It is not clear, a priori, that the boundary of Q is non degenerate, so we must prove that before proceeding.

Applying Lemma 4.6 to H, we see that there are no repeated edges in the boundary of Q. Since ∂Q has 4 edges and G is $(2,2)$-tight, it follows that ∂Q must have at least 3 vertices. Thus the only possibility for a degenerate boundary is that one vertex is repeated and that ∂Q has underlying graph isomorphic to $\mathfrak{B}$. It follows that $\overline{Q}$, the topological closure of the open disc Q in Σ, is homotopy equivalent to an embedded loop. To see this, observe that $\overline{Q}$ can be contracted onto the simple loop formed by the diagonal of Q that joins the repeated vertex in ∂Q to itself. Thus a sufficiently small neighbourhood of $\overline{Q}$ is homeomorphic to an open annulus. It follows that ∂Q is either inessential or annular. By Lemma 6.5 it cannot be annular. On the other hand, if ∂Q is inessential then, since $\mathfrak{B}$ contains a copy of C_2, this contradicts Lemma 6.1. Thus we see that in fact Q must have a non degenerate boundary.

By Lemma 6.4 this means that Q is also a face of G and so there are blockers H_1, H_2, in G, as described by Lemma 4.7. Now consider the Σ-graph $K = H_1 \cup H_2 \cup \partial Q$. This is $(2,2)$-tight, so, by Lemma 3.2, $K \cap H$ is also $(2,2)$-tight. Now, $K \cap H = (H_1 \cap H) \cup (H_2 \cap H) \cup \partial Q$. Using (1), $H_1 \cap H_2 = \emptyset$, $V(H_1 \cap H \cap \partial Q) = \{v_1, v_3\}$, $E(H_1 \cap H \cap \partial Q) = \emptyset$, $V(H_2 \cap H \cap \partial Q) = \{v_2, v_4\}$ and $E(H_2 \cap H \cap \partial Q) = \emptyset$, we have

$$\begin{aligned} 2 &= \gamma(K \cap H) \\ &= \gamma(\partial Q) + \gamma(H_1 \cap H) + \gamma(H_2 \cap H) - \gamma(H_1 \cap H \cap \partial Q) - \gamma(H_2 \cap H \cap \partial Q) \\ &= 4 + \gamma(H_1 \cap H) + \gamma(H_2 \cap H) - 4 - 4. \end{aligned}$$

Thus $\gamma(H_1 \cap H) + \gamma(H_2 \cap H) = 6$. If $\gamma(H_1 \cap H) = 2$ then $(H_1 \cap H) \cup \{v_2\} \cup \{e_1, e_2\}$ would be a type 1 blocker for the contraction G_{Q,v_1,v_3}, contradicting Lemma 4.7. So $\gamma(H_1 \cap H) \geq 3$ and similarly $\gamma(H_2 \cap H) \geq 3$. It follows that $\gamma(H_1 \cap H) = \gamma(H_2 \cap H) = 3$ and that $H_1 \cap H$ and $H_2 \cap H$ are blockers for the contractions H_{Q,v_1,v_3} and H_{Q,v_1,v_3} respectively. Thus both possible contractions of Q are blocked in H as required. □

For example, suppose that Γ is the simple $(2,2)$-tight graph obtained by adding a vertex of degree two to K_4. Is it possible to embed Γ into the torus to create an irreducible torus graph? If so, how many non isomorphic embeddings exist? There are several possible embeddings of Γ to consider, however we can significantly narrow the search space by observing that since K_4 is tight, by Theorem 6.6, any irreducible embedding of Γ must extend an irreducible embedding of K_4. It is not difficult to show that, up to isomorphism there is a unique irreducible embedding of K_4 in the torus (see figure 12). Thus any irreducible embedding of Γ must restrict to this embedding of K_4. Using this observation, it is not difficult to show that, up to isomorphism there are exactly two distinct irreducible torus embeddings of Γ.

7. Irreducible torus graphs

Let $\mathbb{T} = S^1 \times S^1$ be the torus. Throughout this section let $G = (\Gamma, \varphi)$ be an irreducible $(2,2)$-tight $\mathbb{T}$-graph. Our goal in this section is to show that there are only finitely many isomorphism classes of such graphs by establishing an upper bound for the number of vertices of G.

In the case that G is not cellular we will see that we can essentially reduce the problem to the sphere or the annulus. If G is cellular then using Theorem 3.3 and $f_2 = f_3 = 0$ we see that G satisfies $f_5 + 2f_6 + 3f_7 + 4f_8 = 4$ and $f_i = 0$ for $i \geq 9$. Now using $|V| - |E| + \sum_{i\geq 0} f_i = \chi(\mathbb{T}) = 0$

and $|E| = 2|V| - 2$ we deduce that $|V| = 2+f_4+f_5+f_6+f_7+f_8 \leq 2+f_4+f_5+2f_6+3f_7+4f_8 = 6 + f_4$. Thus the problem reduces to establishing a bound for the number of quadrilateral faces that an irreducible $\mathbb{T}$-graph can have.

First we deal with the non cellular case.

Lemma 7.1. *Suppose that $G = (\Gamma, \varphi)$ is an irreducible $(2, 2)$-tight $\mathbb{T}$-graph. If G is not cellular then Γ is either isomorphic to K_1 or to C_2. Furthermore, in the latter case, G is annular.*

Proof. Since Γ is connected it is clear G has a single non cellular face. By cutting along a non separating loop in this face we obtain an $\mathbb{A}$-graph $\hat{G}$. Observe that any face of $\hat{G}$ that is not also a face of G is non cellular. It follows that $\hat{G}$ is an irreducible $\mathbb{A}$-graph. Now the conclusion follows from Theorem 5.4. □

In order to avoid repeating long lists of hypotheses in the next few statements, we fix some notational conventions as follows.

Assumption 1. Suppose that G is cellular, irreducible $(2, 2)$-tight $\mathbb{T}$-graph. Let Q be a quadrilateral face of G with boundary walk $v_1, e_1, v_2, e_2, v_3, e_3, v_4, e_4, v_1$. Let H_1, respectively H_2, be blockers for the contractions G_{Q,v_1,v_3}, respectively G_{Q,v_2,v_4} and let α_1 and α_2 be simple loops intersecting transversely at one point whose construction is described in Section 4.

Lemma 7.2. *Given Assumption 1 at least one of H_1 or H_2 is an inessential $\mathbb{T}$-subgraph of G.*

Proof. Suppose that both are essential. Then there are non separating simple loops β_1 contained in H_1 and β_2 contained in H_2. Now $H_1 \cap H_2 = \emptyset$, so $i(\beta_1, \beta_2) = 0$. However, it is also clear that $i(\alpha_1, \beta_2) = i(\alpha_2, \beta_1) = 0$. As pointed out in Section 4.3 this implies that $i(\alpha_1, \alpha_2) = 0$, contradicting the fact that these curves intersect transversely at one point. □

We collect some more notational conventions for later use.

Assumption 2. The blocker H_1 is an inessential (and so by Lemma 6.3 it's underlying graph is K_2). The blocker H_2 is maximal with respect to inclusion among all blockers for the contraction G_{Q,v_2,v_4}. Furthermore J is the face of H_2 that contains v_1, v_3.

See Figure 6 for an illustration of these Assumptions 1 and 2 in the case where H_2 is an essential blocker.

Lemma 7.3. *Given Assumptions 1 and 2, then any face of H_2 that is not J is also a face of G.*

Proof. Suppose that $F \neq J$ is a face of H_2. Then $\gamma(H_2 \cup \text{int}_G(F)) \leq \gamma(H_2) = 3$, by Theorem 3.4. Also $v_1, v_3 \notin \text{int}_G(F)$, since $F \neq J$. It follows that $H_2 \cup \text{int}_G(F)$ is a blocker for G_{Q,v_2,v_4} and so by the maximality of H_2, $\text{int}_G(F) \subset H_2$ as required. □

Next we want to examine the structure of H_2. It turns out that there are exactly ten distinct possibilities. If H_2 is inessential then, by Lemma 6.3 it has graph K_2 (Figure 7). On the other hand, if H_2 is essential we have the following.

Lemma 7.4. *Given Assumptions 1 and 2 and also assuming that H_2 is essential, then H_2 is isomorphic as a $\mathbb{T}$-graph to one of the nine $\mathbb{T}$-graphs shown in Figures 8 and 9.*

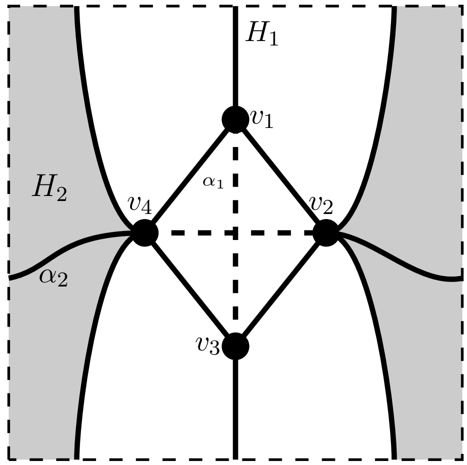

FIGURE 6. An illustration of the notational conventions in Assumptions 1 and 2. The shaded region represents an essential blocker H_2.

Proof. Since $\gamma(H_2) = 3$, it is connected by Lemma 3.1. Let K be the $\mathbb{S}$-graph obtained by cutting and capping H_2 along α_1 (which is a non separating loop in J). Clearly K has two exceptional faces J^+ and J^- such that all other faces of K are faces of G (using Lemma 7.3). Now, since J^+ and J^- are the only faces of K that could have degree less than 4, Theorem 3.3 implies that K satisfies

$$2f_2 + f_3 = 2 + f_5 + 2f_6 \tag{3}$$

and $f_i = 0$ for $i \geq 7$. There are two cases to consider.

(a) There is no quadrilateral face of G in H_2. There are various subcases:

- (a1) $|J^+| = |J^-| = 2$. Then, from Equation 3 we get $f_5 + 2f_6 = 2$. So either $f_5 = 0$ and $f_6 = 1$ and we have the example shown in Figure 8 (a), or, $f_5 = 2$ and $f_6 = 0$ and we have one of the examples shown in Figure 8 (b) or (c).
- (a2) $|J^+| = 2$ and $|J^-| = 3$. Then we have $f_5 = 1$. There is one possibility: Figure 8 (d).
- (a3) $|J^+| = |J^-| = 3$. In this case, Equation 3 implies that J^+ and J^- are the only faces of K. So we have the example shown in Figure 8 (e).
- (a4) $|J^+| = 2$ and $|J^-| = 4$. In this case, Equation 3 implies that J^+ and J^- are the only faces of K and we have the example shown in Figure 8 (f).

(b) There is some quadrilateral face of G in H_2. This case requires a little more effort as we must first establish that there is no more than one such face. Let $G' = \partial Q \cup H_1 \cup H_2$. Clearly G' is $(2,2)$-tight and so by Theorem 6.6 it is also irreducible.

Suppose that R is a quadrilateral face of G, with boundary vertices w_1, w_2, w_3, w_4, that is contained in H_2 (and so is also a face of G'). By Lemma 7.2 we know that there is a blocker for one of the contractions of R in G' whose graph is K_2. Without loss of generality assume that a blocker L_1 for the contraction G'_{R,w_1,w_3} has graph K_2. Now we claim that $L_1 \subset H_2$. If not then it is clear that L_1 must intersect H_1. Since the vertices of L_1 are both in H_2 this contradicts $H_1 \cap H_2 = \emptyset$, thus establishing our claim.

Now consider a maximal blocker, L_2, for the contraction G'_{R,w_2,w_4}. By (1) we have

$$\begin{aligned} 3 &= \gamma(L_2) \\ &= \gamma(L_2 \cap (\partial Q \cup H_1)) + \gamma(L_2 \cap H_2) - \gamma(L_2 \cap (\partial Q \cup H_1) \cap H_2) \\ &= \gamma(L_2 \cap (\partial Q \cup H_1)) + \gamma(L_2 \cap H_2) - \gamma(L_2 \cap \{v_2, v_4\}) \end{aligned}$$

Now it is clear that $\{v_2, v_4\} \subset L_2$ since L_2 is connected, so we have

$$\gamma(L_2 \cap H_2) = 7 - \gamma(L_2 \cap (\partial Q \cup H_1)) \tag{4}$$

Furthermore, it is also clear that L_1 separates v_2 from v_4 in H_2, so $L_2 \cap H_2$ has at least two components, and hence $\gamma(L_2 \cap H_2) \geq 4$. Also $L_2 \cap (\partial Q \cup H_1)$ is a $\mathbb{T}$-subgraph of $\partial Q \cup H_1$ that contains the vertices v_2, v_4. It follows easily that $\gamma(L_2 \cap (\partial Q \cup H_1)) \geq 3$ with equality only if $L_2 \cap (\partial Q \cup H_1) = \partial Q \cup H_1$. Therefore the only way that (4) can be satisfied is that $\partial Q \cup H_1 \subset L_2$ and $L_2 \cap H_2$ has exactly two components $X_2 \ni v_2$ and $X_4 \ni v_4$ such that $\gamma(X_2) = \gamma(X_4) = 2$. In particular it follows from Theorem 6.6 and Lemma 7.1 that the underlying graph of X_2, and also of X_4, is isomorphic to K_1 or C_2. Now since L_1 also separates w_2 and w_4 in H_1 we can, without loss of generality, assume that $v_2, w_2 \in X_2$ and $v_4, w_4 \in X_4$.

Let Z_2, respectively Z_4, be the maximal $(2,2)$-tight $\mathbb{T}$-subgraph of H_2 that contains v_2, respectively v_4. By Lemma 3.2 we see that $X_2 \subset Z_2$ and $X_4 \subset Z_4$. Furthermore we see that since Z_2 and Z_4 are both disjoint from α_1, they are either annular or inessential. By Lemma 7.1, Z_2 has graph K_1 (inessential case) or C_2 (annular case). Similar comments apply to Z_4. Now the argument in the paragraph above shows that every quadrilateral face of H_2 has a boundary vertex in Z_2 and a diagonally opposite vertex in Z_4. It follows easily that there is at most one such quadrilateral face in H_2.

Now we can argue as in case (a) but with the proviso that there is exactly one quadrilateral face, R, of H_2 that is also a face of G. We observe that there is a cycle of length 3 in H_2 (formed by two edges of ∂R and the inessential blocker for R) and so also in K. It is not hard to see that it follows that K must have at least two faces of odd degree: at least one on either 'side' of the cycle of length 3. We find the following subcases.

- (b1) $|J^+| = |J^-| = 2$. From Equation (3) we have $f_5 + 2f_6 = 2$. Since K has some face of odd degree we can rule out the possibility $f_5 = 0, f_6 = 1$. Therefore $f_5 = 2$ and $f_6 = 0$. There is only one possibility for H_2: Figure 9 (a).
- (b2) $|J^+| = 2$ and $|J^-| = 3$. Then, as in case (a) we have $f_5 = 1$ and there is one possibility: Figure 9 (b).
- (b3) $|J^+| = |J^-| = 3$. In this case, Equation 3 implies that R, J^+ and J^- are the only faces of K: Figure 9 (c).
- (b4) $|J^+| = 4$ and $|J^-| = 2$. Since K must have at least two faces of odd degree, Theorem 3.3 would imply that there is a triangle or digon in K that is also a face of G, contradicting its irreducibility. Thus this subcase cannot arise.

□

It remains to rule out the possibility of a quadrilateral face that is neither Q nor a face of H_2. In fact we can prove something a little more general than that.

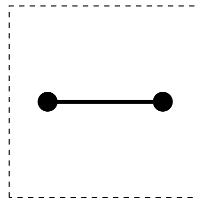

FIGURE 7. The unique inessential blocker

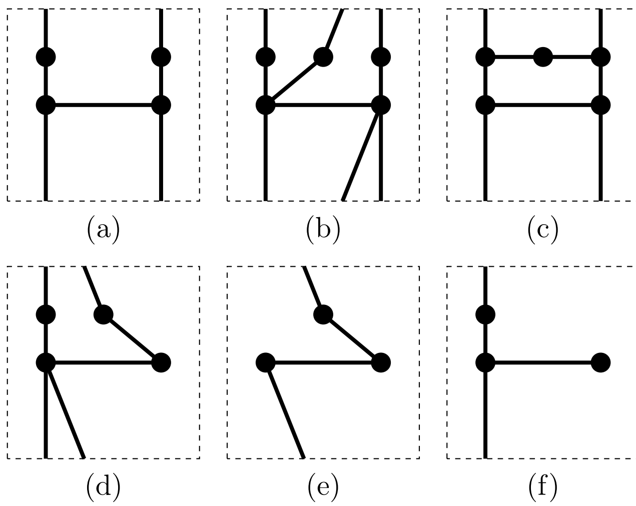

FIGURE 8. Essential blockers with no quadrilateral face

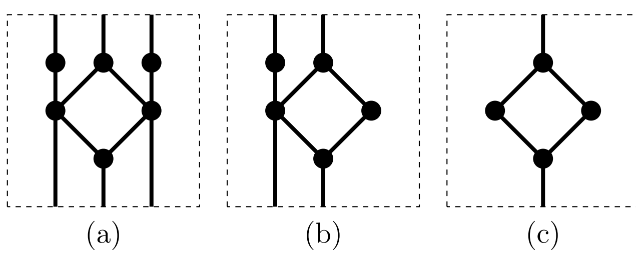

FIGURE 9. Essential blockers with a quadrilateral face

Lemma 7.5. *Suppose that G is an irreducible $(2,2)$-tight $\mathbb{T}$-graph. Let K be a $(2,2)$-tight $\mathbb{T}$-subgraph of G and suppose that F is a cellular face of K. There is no quadrilateral face of G properly contained within F.*

Proof. Suppose that R, with vertices w_1, w_2, w_3, w_4, is a quadrilateral face of G properly contained within F and let B_1 and B_2 be blockers for contractions of R in G. If $B_1 \subset F$, then since F is cellular, B_1 would separate w_2 from w_4 which contradicts $B_1 \cap B_2 = \emptyset$. Therefore B_1 is not contained in F or equivalently, since B_1 is connected, $B_1 \cap K \neq \emptyset$. Similarly $B_2 \cap K \neq \emptyset$.

Now, let $M = \partial R \cup B_1 \cup B_2$ and observe that M is $(2,2)$-tight and therefore, by Lemma 3.2, $M \cap K$ is also $(2,2)$-tight. Now it is clear that $M \cap K = (B_1 \cap K) \cup (B_2 \cap K) \cup E(\partial R \cap K)$. Therefore

$$2 = \gamma(M \cap K) = \gamma(B_1 \cap K) + \gamma(B_2 \cap K) - |E(\partial R \cap K)| \tag{5}$$

Now, we observe that $|E(\partial R \cap K)| \in \{0, 1, 2, 4\}$ since K is an induced $\mathbb{T}$-subgraph of G. If $|E(\partial R \cap K)| = 4$ then clearly R must be a face of K which contradicts our assumption that R is properly contained within F. On the other hand if $|E(\partial R \cap K)| \leq 1$, then (5) yields $\gamma(B_1 \cap K) + \gamma(B_2 \cap K) \leq 3$ which contradicts the fact that both $B_1 \cap K$ and $B_2 \cap K$ are nonempty. Finally if $|E(\partial R \cap K)| = 2$ then it is clear that K contains exactly three of the vertices w_1, w_2, w_3, w_4. However in this case (5) implies that $\gamma(B_1 \cap K) = \gamma(B_2 \cap K) = 2$. It follows that K contains at most one of the vertices w_1, w_3, otherwise $(B_1 \cap K) \cup \{w_2\}$ would span a type 1 blocker for G_{R,w_1,w_3}, contradicting Lemma 4.7. Similarly K contains at most one of the vertices w_2, w_4. Thus K contains at most two of the vertices w_1, w_2, w_3, w_4 yielding the required contradiction. □

Finally we have our main theorem about torus graphs.

Theorem 7.6. *Suppose that G is an irreducible $(2,2)$-tight $\mathbb{T}$-graph. Then G has at most two quadrilateral faces.*

Proof. Suppose, as above, that Q is a quadrilateral face of G, with maximal blockers H_1 and H_2. Also by Lemma 7.2 we may assume that H_1 is inessential. By Lemma 7.4 there is at most one other quadrilateral face of G contained among faces of H_2. Now let $K = \partial Q \cup H_1 \cup H_2$. Clearly K is a $(2,2)$-tight $\mathbb{T}$-subgraph of G. Now we consider the faces of K that are not also faces of G. If H_2 is also inessential then, using Lemma 6.3, there is at most one such face and that face is cellular of degree 8. If H_2 is essential then there at most two such faces and each such face is cellular and has degree at least 5. So by Lemma 7.5, there is no quadrilateral face of G that is not also a face of K. □

Corollary 7.7. *There are finitely many distinct isomorphism classes of irreducible $(2,2)$-tight $\mathbb{T}$-graphs. In particular any such irreducible $\mathbb{T}$-graph has at most eight vertices.*

Proof. We may as well assume that G is cellular, since in the non cellular case we know that G has at most two vertices. Since $\gamma(G) = 2$ we have $|V| = 1 + \frac{1}{4}\sum i f_i$, so we must maximise $\sum i f_i$. Since G is irreducible, $f_i = 0$ for $i = 0, 1, 2, 3$ and $f_4 \leq 2$. From Theorem 3.3 we have $f_5 + 2f_6 + 3f_7 + 4f_8 = 4$ and $f_i = 0$ for $i \geq 9$. Clearly the maximum value for $\sum i f_i$ is attained by having $f_4 = 2$, $f_5 = 4$ and $f_i = 0$ for $i \neq 4, 5$. In that case $|V| = 8$. Now there are finitely many isomorphism classes of $(2,2)$-tight graphs with at most eight vertices. Moreover, for each such graph, there are finitely many isomorphism classes of torus graphs with that underlying graph. □

7.1. **Identifying irreducibles.** Given Corollary 7.7, a naive algorithm to find all the irreducibles mentioned therein would be

1. Find all $(2,2)$-tight graphs with at most 8 vertices.
2. For each such graph, find all isomorphism classes of torus embeddings.
3. Eliminate all embeddings that are not irreducible.

It is impractical to carry out this procedure without the assistance of a computer as step (1) will already yield many thousands of distinct graphs, each of which could have many different torus embeddings.

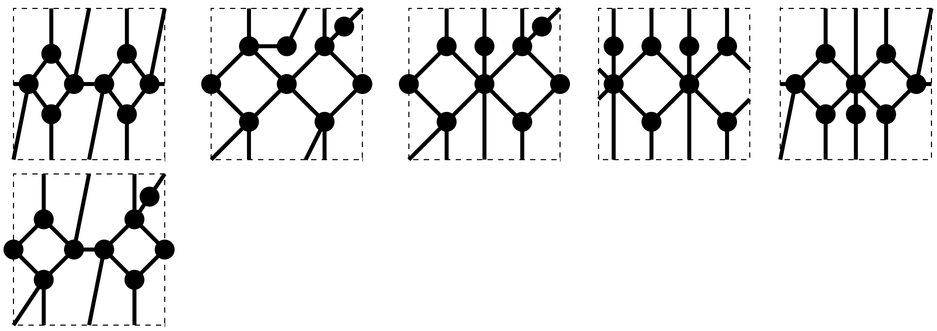

Figure 10. Irreducible torus graphs with eight vertices

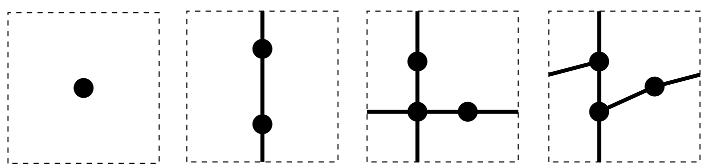

Figure 11. Irreducible torus graphs with at most three vertices

However, since we have a lot of structural information about irreducibles, we can narrow the search space significantly. For example, it is clear from the proof of Corollary 7.7 that any irreducible with 8 vertices must have 2 quadrilateral faces, 4 faces of degree 5 and no other faces. Moreover, we know that each quadrilateral face has one inessential blocker and one other blocker which must be one of the 10 torus graphs described in Figures 7, 8 and 9. It is not too difficult to deduce that any 8 vertex irreducible must be isomorphic to one of the examples shown in Figure 10.

Similarly for torus graphs with at most 4 vertices there are relatively few possibilities for the underlying graph: 13 in total. Now, using Lemmas 6.1, 6.4 and 6.5 we can easily deduce that an irreducible with at most 4 vertices is isomorphic to one of the examples shown in Figures 11 or 12. For the cases of 5, 6 and 7 vertices this naive approach yields a relatively manageable problem in computational graph theory. We have used the computer algebra system SageMath [19] to automate much of the search process in these cases. The interested reader can find full details of the search algorithm and its implementation at [5]. As a result of this computation we have the following.

Theorem 7.8. *There are 116 distinct isomorphism classes of* $(2,2)$*-tight irreducible torus graphs.* □

At the end of the next section, we will describe an alternative inductive construction for $(2,2)$-tight irreducible torus graphs that has fewer irreducibles but requires an additional inductive operation. See Theorem 8.4 below.

8. Application: contacts of circular arcs

In this section we describe an application to the study of contact graphs. The foundational result in this area is the well known Koebe-Andreev-Thurston Circle Packing Theorem ([13]). More recently, contact graphs for many different classes of geometric objects have been studied, with various restrictions placed on the allowed contacts. See for example [1,8,9]

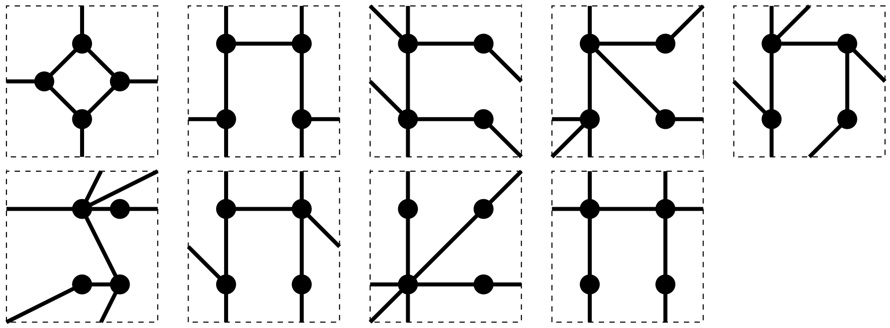

FIGURE 12. Irreducible torus graphs with four vertices

We consider contact graphs arising from certain families of curves in surfaces of constant curvature. We begin by giving a model for a general class of contact problems and then specialise to a case of particular interest.

Let $\alpha : [0,1] \to \Sigma$ be a curve. We say that α is non self-overlapping if it is injective on the open interval $(0,1)$. Now suppose that $\alpha, \beta : [0,1] \to \Sigma$ are distinct curves in Σ. We say that α and β are non overlapping if $\alpha((0,1)) \cap \beta((0,1)) = \emptyset$. Let $\mathcal{C}$ be a collection of curves in Σ having the following properties

- Every $\alpha \in \mathcal{C}$ is non self-overlapping.
- For every distinct $\alpha, \beta \in \mathcal{C}$, α and β are non overlapping.

We want to construct a combinatorial object that describes the contact properties of such a collection. In order to do this we impose some further non degeneracy conditions on $\mathcal{C}$ as follows.

- $\alpha(0) \neq \alpha(1)$ for every $\alpha \in \mathcal{C}$
- For every distinct $\alpha, \beta \in \mathcal{C}$, $\{\alpha(0), \alpha(1)\} \cap \{\beta(0), \beta(1)\}$ is empty.

In other words, we allow the end of one curve to touch another curve (or to touch itself), but the point that it touches cannot be an endpoint of that curve. We say that $\mathcal{C}$ is a non degenerate collection of non overlapping curves. Note that if a collection fails the non degeneracy conditions, it can typically be made non degenerate by an arbitrarily small perturbation. A contact of $\mathcal{C}$ is a quadruple (α, β, x, y) where $\alpha, \beta \in \mathcal{C}$, $x \in \{0,1\}$, $y \in (0,1)$ and $\alpha(x) = \beta(y)$.

Now we can define a graph $\Gamma_{\mathcal{C}}$ as follows. The vertex set is $\mathcal{C}$ and the edge set is $\mathcal{T}$, the set of contacts of $\mathcal{C}$. We define the incidence functions $s, t : \mathcal{T} \to \mathcal{C}$ by $s(\alpha, \beta, x, y) = \alpha$ and $t(\alpha, \beta, x, y) = \beta$. The quadruple $(\mathcal{C}, \mathcal{T}, s, t)$ is a directed multigraph and we let $\Gamma_{\mathcal{C}}$ be the graph obtained by forgetting the edge orientations. We can construct an embedding $|\Gamma_{\mathcal{C}}| \to \Sigma$ as follows. For $\beta \in \mathcal{C}$, suppose that $t^{-1}(\beta) = \{(\alpha_1, \beta, x_1, y_1), \cdots, (\alpha_k, \beta, x_k, y_k)\}$. Let J_β be a nonempty closed subinterval of $[0,1]$ with the following properties.

(1) $\{y_1, \cdots, y_k\} \subset J_\beta$.
(2) $0 \in J_\beta$ if and only if there is no contact $(\beta, \gamma, 0, y)$ in $\mathcal{T}$.
(3) $1 \in J_\beta$ if and only if there is no contact $(\beta, \gamma, 1, y)$ in $\mathcal{T}$.

In other words J_β is a subinterval that covers all the 'points of contact' in β together with any endpoints of β that do not touch a curve. Now we observe that the restriction $\beta|_{J_\beta}$ is injective and therefore is a homeomorphism onto its image. So it follows from the Jordan-Schoenflies Theorem that $\Sigma/\beta(J_\beta)$ is homeomorphic to Σ. Furthermore, since $\beta(J_\beta) \cap \delta(J_\delta) = \emptyset$ for $\beta \neq \delta$ it follows that Σ is homeomorphic to $\Sigma/\sim$ where $\sim$ is the equivalence relation that collapses

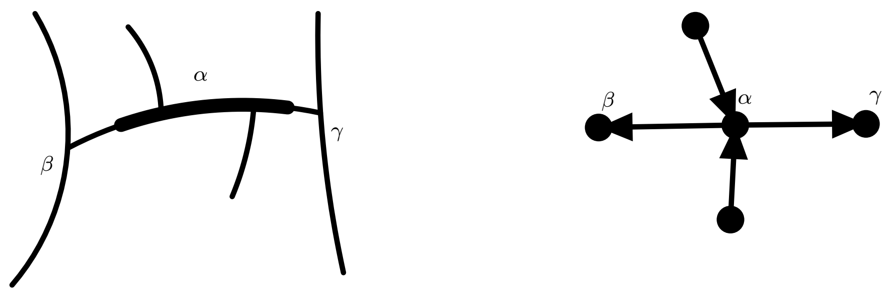

FIGURE 13. The construction of the contact graph associated to a collection of curves. On the left we have a collection of curves. The bold section of α represents $\alpha(J_\alpha)$. On the right is the corresponding graph with edge orientations as indicated.

each $\beta(J_\beta)$ to a point, for all $\beta \in \mathcal{C}$. Using this homeomorphism we construct an embedding $\varphi : |\Gamma_\mathcal{C}| \to \Sigma/\sim$ by mapping each vertex of $\Gamma_\mathcal{C}$ (i.e. element of $\mathcal{C}$) to the corresponding point of $\Sigma/\sim$. Since an edge of $\Gamma_\mathcal{C}$ is a contact (α, β, x, y), we can construct the corresponding edge embedding by using the restriction of α to the component of $[0, 1] - J_\alpha$ that contains x. The contact graph of $\mathcal{C}$ is defined to be the $(\Sigma/\sim)$-graph $(\Gamma_\mathcal{C}, \varphi)$: see Figure 13 for an illustration of this construction. Note that the term contact graph is used in a variety of different ways in the literature depending on the context. In the remainder of this paper we shall always use the term contact graph to mean the surface graph described above.

We are interested in the recognition problem for contact graphs: can we find necessary and/or sufficient conditions for a surface graph to be the contact graph of a collection of curves? Typically we are looking for conditions for which there is an efficient decision algorithm. As noted in the introduction, there are efficient algorithms for deciding whether or not a given graph is $(2, l)$-sparse. See [14] and [10] for details.

Hliněný ([12]) has shown that a simple plane graph is the contact graph of a generic collection of plane curves if and only if it is $(2, 0)$-sparse. We note that he uses a slightly different notion of generic than we have and he states the result in terms of the intersection graph of the collection of curves. However it is easy to see that the result and the method of proof adapt easily to the model of contact graphs described above and to arbitrary surfaces. We do not need this result but nevertheless include a precise statement to provide some context for our later result.

Theorem 8.1. *Let G be a Σ-graph. Then $G \cong G_\mathcal{C}$ for some collection of curves $\mathcal{C}$ in Σ if and only if G is $(2, 0)$-sparse.* □

It is worth noting here that graphs studied in [12] and elsewhere are typically the intersection graphs of a collection of curves. While this is usually appropriate in the plane, for non simply connected surfaces we propose that it is more natural to consider the contact graph as defined above. We observe that the intersection graph is obtained from the surface graph $G_\mathcal{L}$ by taking the underlying graph, deleting all loops and replacing any sets of parallel edges by a single edge.

Now we suppose that Σ is also equipped with a metric of constant curvature. In this context we can distinguish many interesting subclasses of curves. For example, a circular arc is a curve of constant curvature and a line segment is a locally geodesic curve. For collections of such curves the recognition question can depend on the embedding of the graph and not just the

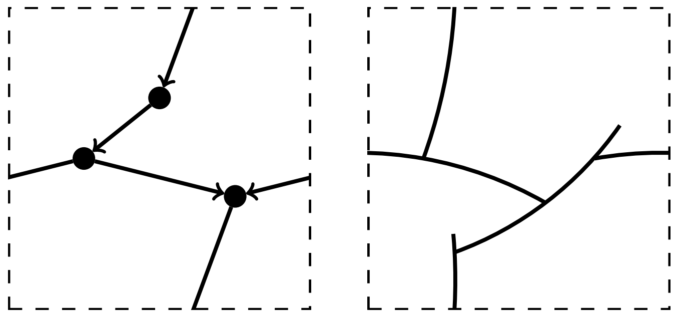

FIGURE 14. A torus graph and a corresponding CCA representation in the flat torus. The orientation of the graph edges is the orientation induced by the CCA representation.

graph itself (in contrast to Theorem 8.1). For example, the graph C_2 has two non isomorphic embeddings in the torus, one essential and one inessential. It is easy to see that the essential embedding results in a surface graph that is the contact graph of a pair of line segments in the flat torus. On the other hand the inessential embedding has no such representation by a pair of line segments.

Given a Σ-graph G and a non degenerate non overlapping collection of circular arcs C such that $G \cong G_{\mathcal{C}}$ we say that $\mathcal{C}$ is a CCA representation of $\mathcal{C}$ (abbreviating Contacts of Circular Arcs). See Figure 14 for an example in the torus. Alam et al. ([1]) have shown that any $(2,2)$-sparse plane graph has a CCA representation in the flat plane. We prove an analogous result for the flat torus.

First we need a lemma to show that every sparse surface graph can be obtained by deleting only edges from a tight surface graph.

Lemma 8.2. *Suppose that Σ is a connected surface, $l \leq 2$ and G is a $(2,l)$-sparse Σ-graph. There is some $(2,l)$-tight Σ-graph H such that $V(H) = V(G)$ and G is a Σ-subgraph of H.*

Proof. Clearly it suffices to show that if $\gamma(G) \geq l+1$ then we can add an edge e within some face of G so that $G \cup \{e\}$ is $(2,l)$-sparse.

So suppose that $\gamma(G) \geq l+1$. Now if G has no tight Σ-subgraph then we can add any edge without violating the sparsity count. So we assume that G has some nonempty tight Σ-subgraph. Let L be a maximal tight Σ-subgraph of G. If $V(L) = V(G)$ then, $\gamma(G) \leq \gamma(L) = l$, which contradicts our assumption, so $V(L)$ is a proper subset of $V(G)$. Since Σ is connected there is some face F of G whose boundary contains vertices $u \in L$ and $v \notin L$. Let e be a new edge that joins u and v through a path in F. We claim that $G \cup \{e\}$ is $(2,l)$-sparse. If not then there must be some tight Σ-subgraph K of G such that $u, v \in K$. But $K \cap L$ is nonempty, so by Lemma 3.2, $K \cup L$ is $(2,l)$-tight. This contradicts the maximality of L. □

Theorem 8.3. *Every $(2,2)$-sparse torus graph admits a CCA representation in the flat torus.*

Proof sketch. First observe that edge deletion is CCA representable: just shorten one of the arcs slightly. So by Lemma 8.2 it suffices to prove the theorem for $(2,2)$-tight torus graphs. To that end we must show that

(a) each irreducible $(2,2)$-tight torus graph has a CCA representation.

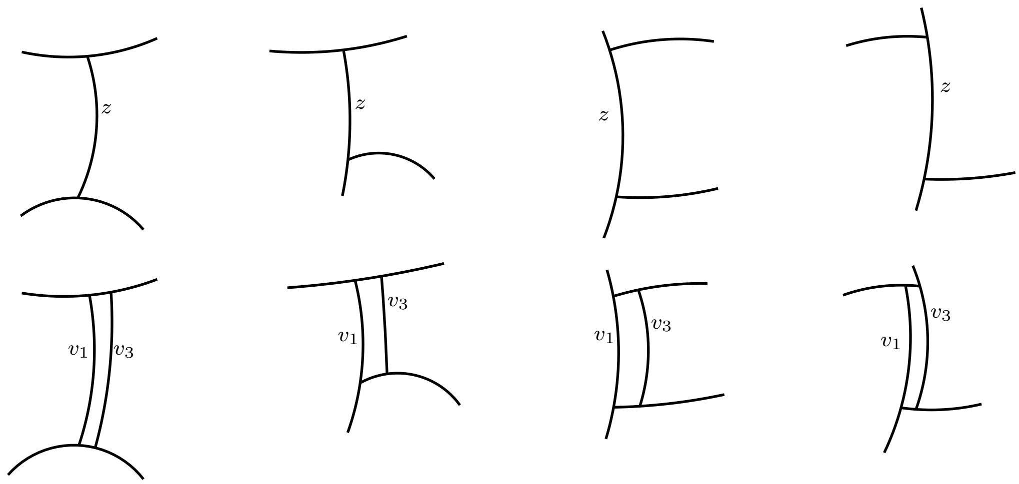

FIGURE 15. CCA representations of quadrilateral splitting moves. The configurations in the top row are CCA representations of a graph consisting of two edges incident to a vertex (i.e circular arc) z and are representative of the various possibilities depending on the orientations of the contacts at either end of the arc z. The configurations in the bottom row are the CCA representations of the graph obtained by applying a quadrilateral split to the two incident edges that replaces z by arcs v_1, v_3. Observe that in each case v_1 and v_3 can be chosen so that they are contained in an arbitrarily small neighbourhood of the arc z. Thus it is clear that any other arc (not shown in the diagram) that might be incident with z in the top configuration can be perturbed so that it is incident with either v_1 or v_3 as appropriate in the bottom configuration.

(b) if $G \to G'$ is a digon, triangle or quadrilateral contraction move and G' has a CCA representation, then G also has a CCA representation. In other words the relevant vertex splitting moves are CCA representable.

For (a) it is possible to give an explicit CCA representation for each of the 116 irreducibles referred to in Theorem 7.8. We will not describe those here but below we shall explain a simple method to make these constructions easily. Full details are given in [20].

For (b), see Figure 8 for an illustration of the various possible CCA representations of a quadrilateral split. It should be apparent from this diagram that any quadrilateral split is CCA representable. The digon and triangle vertex splits are also representable and indeed have already been dealt with in the plane context in [1]. We observe that the constructions described there work equally well for torus graphs.

In order to construct the CCA representations for the 116 irreducible torus graphs mentioned in Theorem 7.8, we can make use of topological Henneberg moves. We remind the reader that a Henneberg vertex addition move is the operation of adding a new vertex to a graph and two edges from that vertex to the existing graph. Note that we allow the two new edges to be parallel. Moreover in the context of surface graphs we insist that the new vertex is placed in some face of the existing graph and the two edges are incident with vertices in the boundary of that face. We refer to such an operation as a topological Henneberg move. Clearly a Henneberg move is the inverse operation to divalent vertex deletion. It is well known (and elementary) that divalent vertex deletions preserve $(2, l)$-sparseness for all l. On the other

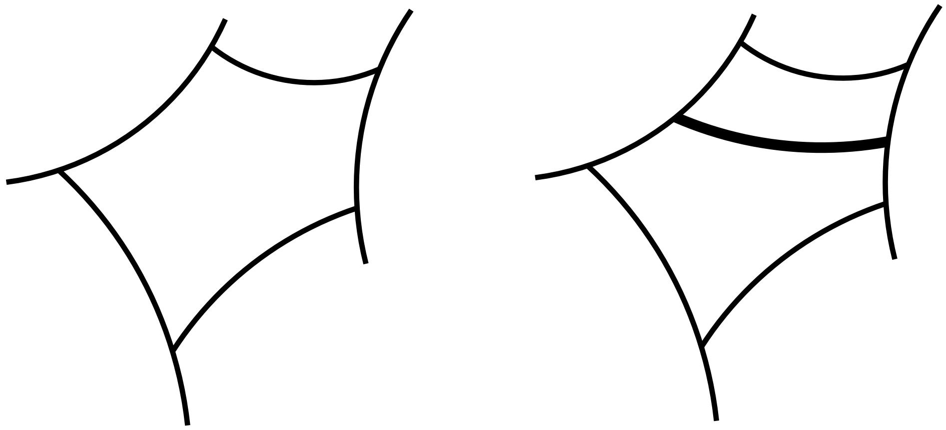

FIGURE 16. A CCA representation of a topological Henneberg move. The bold arc on the right represents the new vertex, that touches two of the initial arcs.

hand Henneberg moves preserve $(2,l)$-sparseness for $l \leq 2$, and for $l = 3$ if we also insist that the new edges are not parallel. [4]

It turns out that there are just 12 irreducibles that have no vertices of degree 2. In Figure 17 we give diagrams of each of these torus graphs and in Figure 18 we give sample CCA representations of each of these in the flat torus. We observe that each of the 116 irreducible graphs can be constructed by a sequence topological Henneberg moves from one of the torus graphs in Figure 17: indeed one easily sees that at most five Henneberg moves are required.

It remains to show that the required topological Henneberg moves are CCA representable. A CCA representation of a topological Henneberg move is illustrated in Figure 16. In general of course, topological Henneberg moves can fail to be CCA representable given a fixed representation of the initial graph. It is easy to construct a CCA representation of a graph that has a highly non convex region and so may not admit the required circular arc to realise a topological Henneberg move. However, it is readily verified that given the CCA representations in Figure 18, it is possible to represent all the necessary Henneberg moves that are required to construct CCA representations of the full set of 116 irreducible graphs. See [20] for complete details of this. □

Finally we observe that allowing divalent vertex additions we have the following inductive construction for $(2,2)$-tight torus graphs.

Theorem 8.4. *If G is a $(2,2)$-tight torus graph then G can be constructed from one of the torus graphs in Figure 17 by a sequence of moves each of which is either a digon split, triangle split, quadrilateral split or a divalent vertex addition.* □

[4]A graph Γ is $(2,3)$-sparse if $\gamma(\Lambda) \geq 3$ for every subgraph Λ of Γ that contains at least one edge.

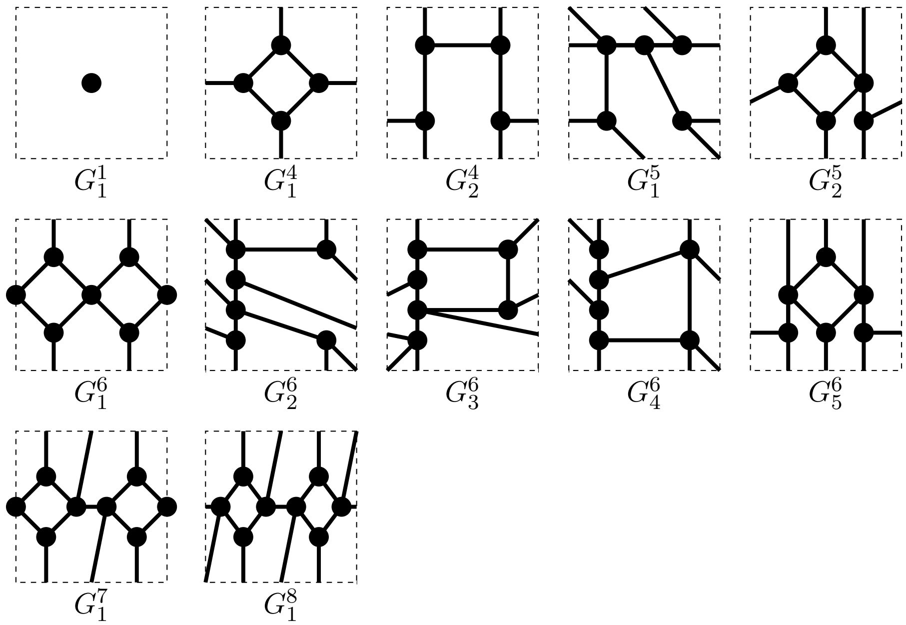

FIGURE 17. The irreducibles that have no vertex of degree 2.

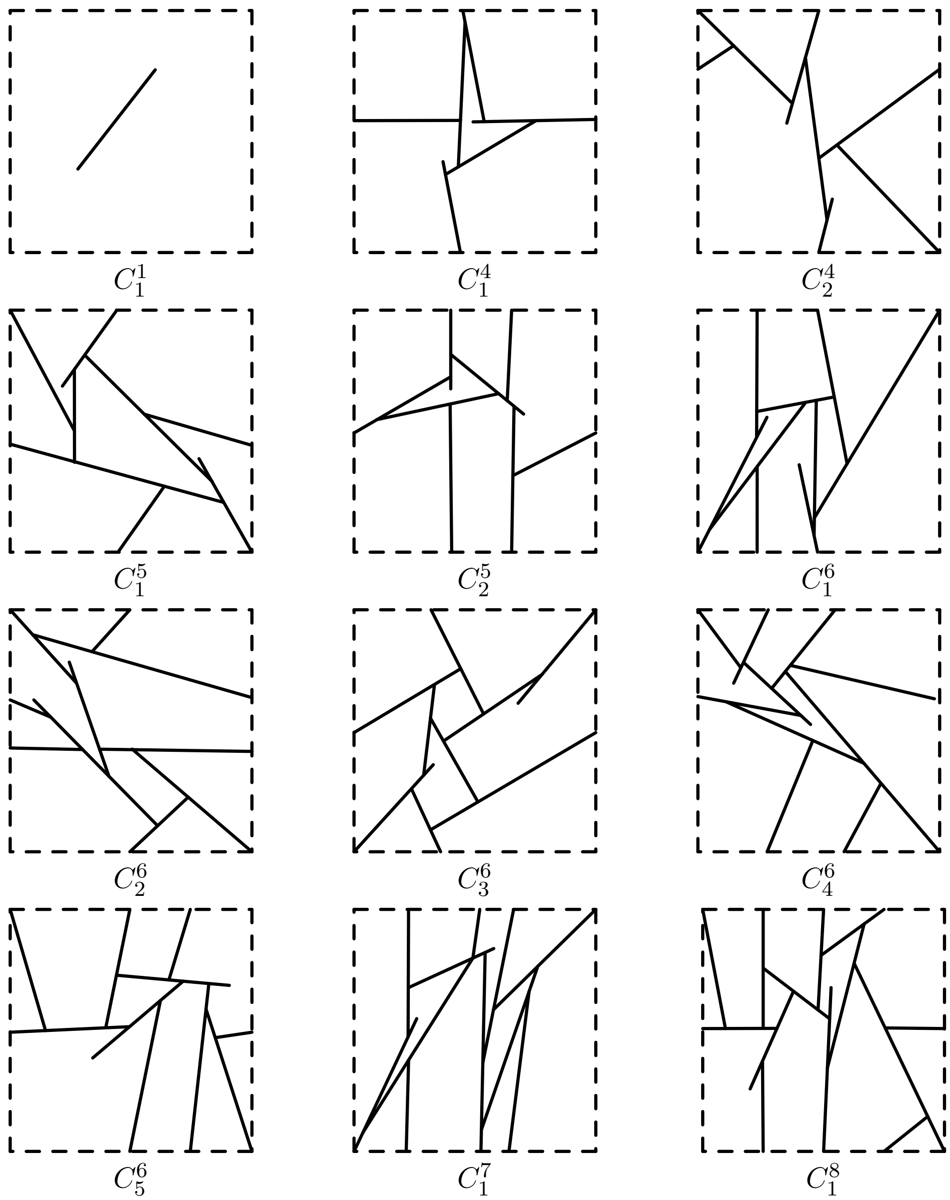

FIGURE 18. CCA representations of the 12 irreducible torus graphs with no vertices of degree two: C_j^i is a CCA representation of G_j^i from Figure 17.

References

[1] Md. Jawaherul Alam, David Eppstein, Michael Kaufmann, Stephen G. Kobourov, Sergey Pupyrev, André Schulz, and Torsten Ueckerdt, *Contact graphs of circular arcs*, Algorithms and data structures, Lecture Notes in Comput. Sci., vol. 9214, Springer, Cham, 2015, pp. 1–13. MR3677547

[2] D. W. Barnette and Allan L. Edelson, *All 2-manifolds have finitely many minimal triangulations*, Israel J. Math. **67** (1989), no. 1, 123–128, DOI 10.1007/BF02764905. MR1021367

[3] James Cruickshank, Derek Kitson, and Stephen C. Power, *The generic rigidity of triangulated spheres with blocks and holes*, J. Combin. Theory Ser. B **122** (2017), 550–577. MR3575219

[4] J. Cruickshank, D. Kitson, and S. C. Power, *The rigidity of a partially triangulated torus*, Proc. Lond. Math. Soc. (3) **118** (2019), no. 5, 1277–1304, DOI 10.1112/plms.12215. MR3946722

[5] James Cruickshank and Qays Shakir, *Software for computing irreducible* $(2, 2)$*-tight torus graphs*, 2021, DOI 10.5281/zenodo.4579832, available at `http://doi.org/10.5281/zenodo.4579832`

[6] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012. MR2850125

[7] Zsolt Fekete, Tibor Jordán, and Walter Whiteley, *An inductive construction for plane Laman graphs via vertex splitting*, Algorithms—ESA 2004, Lecture Notes in Comput. Sci., vol. 3221, Springer, Berlin, 2004, pp. 299–310. MR2165974

[8] Hubert de Fraysseix and Patrice Ossona de Mendez, *Representations by contact and intersection of segments*, Algorithmica **47** (2007), no. 4, 453–463, DOI 10.1007/s00453-006-0157-x. MR2304990

[9] Daniel Gonçalves, Lucas Isenmann, and Claire Pennarun, *Planar graphs as* ⌞*-intersection or* ⌞*-contact graphs*, Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 2018, pp. 172–184, DOI 10.1137/1.9781611975031.12. MR3775800

[10] Donald J. Jacobs and Bruce Hendrickson, *An algorithm for two-dimensional rigidity percolation: the pebble game*, J. Comput. Phys. **137** (1997), no. 2, 346–365, DOI 10.1006/jcph.1997.5809. MR1481894

[11] Ruth Haas, David Orden, Günter Rote, Francisco Santos, Brigitte Servatius, Herman Servatius, Diane Souvaine, Ileana Streinu, and Walter Whiteley, *Planar minimally rigid graphs and pseudo-triangulations*, Comput. Geom. **31** (2005), no. 1-2, 31–61, DOI 10.1016/j.comgeo.2004.07.003. MR2131802

[12] Petr Hliněný, *Classes and recognition of curve contact graphs*, J. Combin. Theory Ser. B **74** (1998), no. 1, 87–103, DOI 10.1006/jctb.1998.1846. MR1644051

[13] P. Koebe, *Kontaktprobleme der konformen Abbildung.*, Ber. Sächs. Akad. Wiss. Leipzig, Math.-phys. (1936), 88:141-164 (German).

[14] Audrey Lee and Ileana Streinu, *Pebble game algorithms and sparse graphs*, Discrete Math. **308** (2008), no. 8, 1425–1437, DOI 10.1016/j.disc.2007.07.104. MR2392060

[15] Bojan Mohar and Carsten Thomassen, *Graphs on surfaces*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 2001. MR1844449

[16] Atsuhiro Nakamoto and Katsuhiro Ota, *Note on irreducible triangulations of surfaces*, J. Graph Theory **20** (1995), no. 2, 227–233, DOI 10.1002/jgt.3190200211. MR1348564

[17] Atsuhiro Nakamoto, *Irreducible quadrangulations of the torus*, J. Combin. Theory Ser. B **67** (1996), no. 2, 183–201, DOI 10.1006/jctb.1996.0040. MR1399674

[18] A. Nixon, J. C. Owen, and S. C. Power, *Rigidity of frameworks supported on surfaces*, SIAM J. Discrete Math. **26** (2012), no. 4, 1733–1757, DOI 10.1137/110848852. MR3022162

[19] The Sage Developers, *Sagemath, the Sage Mathematics Software System (Version 8.1)*, 2017, available at `https://www.sagemath.org`

[20] Qays Shakir, *Sparse topological graphs*, Ph.D. Thesis, National University of Ireland Galway, 2020.

9. Declarations

9.1. **Funding.** The second and third authors were supported by the UK Engineering and Physical Research Council (grant number EP/P01108X/1)

The fourth author was supported by the Iraqi Ministry of Higher Education and Scientific Research and by Middle Technical University, Baghdad.

9.2. **Conflicts of interest/Competing interests.** Not applicable

9.3. **Availability of data and material.** Not applicable

9.4. **Code availability.** All relevant code is freely available at https://zenodo.org/record/4579832

School of Mathematics, Statistics and Applied Mathematics, National University of Ireland Galway, Ireland.
Email address: james.cruickshank@nuigalway.ie

Department of Mathematics and Computer Studies, Mary Immaculate College, Thurles, Ireland.
Email address: derek.kitson@mic.ul.ie

Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, U.K.
Email address: s.power@lancaster.ac.uk

Middle Technical University, Baghdad, Iraq
Email address: qays.shakir@gmail.com