

CONSTRUCTING ISOSTATIC FRAMEWORKS FOR THE ℓ^∞ PLANE

KATIE CLINCH AND DEREK KITSON

ABSTRACT. We use a new coloured multi-graph constructive method to prove that every 2-tree decomposition can be realised in the plane as a bar-joint framework which is minimally rigid (isostatic) with respect to ℓ^1 or ℓ^∞ distance constraints. We show how to adapt this technique to incorporate symmetry and indicate several related open problems on rigidity, redundant rigidity and forced symmetric rigidity in normed spaces.

1. INTRODUCTION

A simple graph $G = (V, E)$ with vertices embedded generically in \mathbb{R}^2 inherits a natural edge-labelling $\kappa : E \rightarrow \{1, 2\}$ whereby an edge (represented by a straight line segment between its embedded vertices) is labelled either 1 or 2 depending on whether the slope m of its affine span satisfies $|m| < 1$ or $|m| > 1$. Simple examples show that not all edge-labellings $\kappa : E \rightarrow \{1, 2\}$ are realisable in this way. Motivated by problems in graph rigidity under ℓ^p distance constraints (see [8, 9, 10]), we are interested in the realisability of d -tree decompositions in \mathbb{R}^d (see Section 2 for the corresponding definition when $d > 2$). A *d-tree decomposition* arises from an edge-labelling $\kappa : E \rightarrow \{1, 2, \dots, d\}$ when the edge sets $\kappa^{-1}(1), \dots, \kappa^{-1}(d)$ are spanning trees in G . In general, multi-graphs which admit a d -tree decomposition are characterised by the conditions $|E| = d(|V| - 1)$ and $|E(H)| \leq d(|V(H)| - 1)$ for each subgraph H (see Nash-Williams [12] and Tutte [16]) and such graphs are said to be *(d, d)-tight*. Constructive characterisations for *(d, d)-tight* graphs and connections to graph rigidity under ℓ^2 distance constraints are discussed in Tay [14], Frank and Szegő [1] and in Graver, Servatius and Servatius [2, §4.9], for example.

It is shown in [8] that, for $p \neq 2$, the structure graph of an isostatic bar-joint framework (see Section 2) in the ℓ^p plane is necessarily *(d, d)-tight*. However, the validity of the converse statement is an open problem for $d \geq 3$ and determining the realisability or otherwise of d -tree decompositions in \mathbb{R}^d would settle this conjecture for the ℓ^∞ context. In this article we present a constructive method for realising coloured 2-tree decompositions in the plane. The method consists of two parts: a multi-graph construction scheme for *(d, d)-tight* graphs which tracks the evolution of d edge-disjoint spanning trees, and in the case $d = 2$, a method of constructing geometric placements for these multi-graphs which accommodates parallel edges. While it is known that *(d, d)-tight* graphs are constructible in terms of multi-graphs, the particular role of spanning trees in these constructions is given prominence here. Moreover, the method of assigning geometric placements

to multi-graphs used here is a new technique which can be adapted for other contexts. The graph construction is presented in Section 3 and the geometric realisations are contained in Section 4. In Section 5, we illustrate the versatility of the technique by adapting it to symmetric 2-tree decompositions. In the concluding section we state some related open problems and indicate connections to other areas of graph rigidity for normed spaces.

This work was initiated at the workshop ‘‘Advances in Combinatorial and Geometric Rigidity’’ which took place at the Banff International Research Station from 12th-17th July 2015.

2. PRELIMINARIES

In this section we introduce terminology, state the main results and provide some background.

2.1. Graph Theory. Let $G = (V, E)$ be a finite, loop-free multi-graph. Let $X \subseteq V$ be a set of vertices. The *neighbourhood* of X in G , $N_G(X)$, is the set of all vertices in $V - X$ which share an edge with some $x \in X$. When $X = \{x\}$ we refer to $N_G(x)$ instead of $N_G(\{x\})$. The subgraph of G induced by X is denoted $G[X]$ and has vertex set X and edge set $E_G(X) = \{uv \in E : u, v \in X\}$. We let $i_G(X) = |E_G(X)|$. When the original graph G is apparent from the context, we omit the subscripts and refer simply to $N(X)$, $E(X)$ and $i(X)$. For $F \subseteq E$, the subgraph of G induced by F is denoted $G[F]$ and has vertex set $V_G(F) = \{v \in V : uv \in F \text{ for some } u \in V\}$, and edge set F . We say that the edge set F spans G when $V_G(F) = V$. Once more, when the graph is apparent, we omit the subscripts. For a vertex $v \in V$, the *degree* of v is the number of edges incident to v , and is denoted by $d_G(v)$ or $d(v)$. Note that for multi-graphs $d_G(v) \geq |N_G(v)|$. The *minimum degree* of G is the minimum of $d(v)$ for all $v \in V$ and is denoted $\delta(G)$.

A *d-tree decomposition* is a tuple $\mathcal{G} = (G; T_1, \dots, T_d)$ where G is a multi-graph which is the edge-disjoint union of spanning trees T_1, \dots, T_d . Note that since G is a multi-graph, it may have multiple edges between a given pair of vertices $u, v \in V(G)$. As such, we say two spanning trees T_i and T_j of G are *edge-disjoint* if whenever they both contain a uv -edge, these edges are distinct in G . Formally, we regard the tuple $\mathcal{K}_1 = (K_1; T_1, \dots, T_d)$, where K_1 is the graph with a single vertex and no edge, and the edge sets of T_1, \dots, T_d are empty, as a *d-tree decomposition*. We denote by \mathcal{G}_d the set of all *d-tree decompositions*. Note that if a multi-graph G admits a *d-tree decomposition* then it is necessarily loop-free and contains at most d copies of any edge.

2.2. Realisations. Let $G = (V, E)$ be a finite simple graph. A *placement* of G in \mathbb{R}^d is an injective map $p : V \rightarrow \mathbb{R}^d$. The pair (G, p) is referred to as a (*bar-joint*) *framework* in \mathbb{R}^d . For each edge $uv \in E$, the pair $\{p(u), p(v)\}$ is said to be *well-positioned* if there exists a unique $k \in \{1, \dots, d\}$ such that

$$\|p(u) - p(v)\|_\infty = |(p(u) - p(v)) \cdot e_k|.$$

Here $(e_i)_{i=1}^d$ is the usual basis for \mathbb{R}^d and $\|x\|_\infty = \max_{1 \leq i \leq d} |x \cdot e_i|$ is the ℓ^∞ norm of a vector $x \in \mathbb{R}^d$. The framework (G, p) is said to be *well-positioned* if $\{p(u), p(v)\}$ is well-positioned for every edge $uv \in E$.




Figure 1. Left and centre: Realisations of the two distinct 2-tree decompositions for the wheel graph W_5 . Right: A placement of W_5 which is not well-positioned.

Each well-positioned framework (G, p) in \mathbb{R}^d inherits an edge-labelling $\kappa_p : E \rightarrow \{1, \dots, d\}$, referred to as the *framework colouring*, where for each edge $uv \in E$,

$$\kappa_p(uv) = \{k\} \text{ if and only if } \|p(u) - p(v)\|_\infty = |(p(u) - p(v)) \cdot e_k|.$$

The set of edges of colour k , $\kappa_p^{-1}(k)$, induces the subgraph $G[\kappa_p^{-1}(k)]$ which is referred to as an induced *monochrome subgraph* of G . A *realisation* for a d -tree decomposition $\mathcal{G} = (G; T_1, \dots, T_d)$, where G is a simple graph, is a framework (G, p) in \mathbb{R}^d with the property that T_1, \dots, T_d are the induced monochrome subgraphs of G .

Example 2.1. Figure 1 illustrates three placements of the wheel graph W_5 in the plane together with the induced framework colourings. Note that W_5 admits exactly two distinct 2-tree decompositions (up to graph isomorphism) and these are realised by the placements in the left and centre of Figure 1. These two placements are well-positioned as each edge affinely spans a line of slope m with $|m| \neq 1$. The rightmost placement is not well-positioned as the edges incident to v_0 affinely span lines of slope ± 1 .

2.3. Graph rigidity. Motivation for considering realisation problems of this type comes from graph rigidity in normed spaces. Consider again a simple graph G . A framework (G, p) is *(locally) rigid* with respect to a given norm on \mathbb{R}^d if every edge-length preserving continuous motion of the vertices is obtained from an isometric motion of the space. The notion of a well-positioned framework (G, p) , introduced above in the context of the ℓ^∞ norm, generalises to the condition that the *rigidity map*,

$$f_G : (\mathbb{R}^d)^V \rightarrow \mathbb{R}^E, \quad (x(v))_{v \in V} \mapsto (\|x(v) - x(w)\|)_{vw \in E},$$

is differentiable at $p = (p(v))_{v \in V}$. In the case of the ℓ^∞ norm on \mathbb{R}^d , *infinitesimal rigidity* is the property that elements which lie in the kernel of the differential $df_G(p)$, also called the *infinitesimal flexes* of the framework, are translational, i.e. take the form $(a, \dots, a) \in (\mathbb{R}^d)^V$ where $a \in \mathbb{R}^d$. It can be shown that, with respect to the ℓ^∞ norm, local rigidity and infinitesimal rigidity are equivalent for well-positioned frameworks (see [6]). In this case we simply refer to a framework as being *rigid*. A framework (G, p) is *minimally rigid* (or *isostatic*) if it is rigid and removing any edge from G results in a framework which is not rigid.

Consider an ℓ^q norm on \mathbb{R}^d with $1 \leq q \leq \infty$ and $q \neq 2$. It is shown in [8] that a necessary condition for a well-positioned framework (G, p) on a

simple graph G to be isostatic with respect to ℓ^q is that G is (d, d) -tight. It is also shown that in the case $d = 2$, a converse statement holds: every $(2, 2)$ -tight simple graph G admits a well-positioned framework (G, p) in the plane which is isostatic for the ℓ^q norm. It is conjectured that this converse statement extends to $d \geq 3$. In the case of the ℓ^∞ norm, the spanning tree characterisation of (d, d) -tight graphs obtained by Nash-Williams [12] and Tutte [16], together with the following result, support this conjecture and provide the link to the realisation problems considered in this article.

Theorem 2.2. [8, Propositions 4.3 & 4.4] *Let G be a simple graph and let (G, p) be a well-positioned framework in $(\mathbb{R}^d, \|\cdot\|_\infty)$. The following statements are equivalent:*

- (a) (G, p) is minimally rigid;
- (b) the monochrome subgraphs induced by the framework colouring κ_p are spanning trees in G .

To summarise, if G is a simple graph which is (d, d) -tight then it admits a d -tree decomposition and so, by the above theorem, to show that G admits a well-positioned minimally rigid framework in $(\mathbb{R}^d, \|\cdot\|_\infty)$ it is sufficient to prove that this d -tree decomposition can be realised in \mathbb{R}^d .

2.4. Inductive constructions. Nash-Williams [12] and Tutte [16] independently characterised the multi-graphs which admit a d -tree decomposition as those which are (d, d) -tight:

Theorem 2.3. *A multi-graph $G = (V, E)$ is expressible as an edge-disjoint union of d spanning trees if and only if*

- (a) $|E| = d|V| - d$, and
- (b) $i(X) \leq d|X| - d$ for all $\emptyset \neq X \subseteq V$.

During his work on the rigidity of body-bar frameworks, Tay [14] found the following inductive construction of (d, d) -tight graphs:

Theorem 2.4. *A multi-graph $G = (V, E)$ is (d, d) -tight if and only if there exists a sequence of graphs*

$$K_1 = G^{(1)} \rightarrow G^{(2)} \rightarrow \dots \rightarrow G^{(n)} = G$$

such that for all $2 \leq i \leq n$, $G^{(i)}$ is obtained from $G^{(i-1)}$ by a d -dimensional j -extension, for some $0 \leq j \leq d - 1$.

A d -dimensional j -extension of a graph $G = (V, E)$ forms a new graph G' by first deleting some set of edges F from G , with $|F| = j$, and then appending a new vertex v to G , incident to $d + j$ new edges, such that $N_{G'}(v) \supseteq V_G(F)$. Note that in rigidity theory it is usually sufficient to only consider $j \leq d - 1$, however in this article we will allow $j \geq d$ when $d = 1$. The inverse of a d -dimensional j -extension is a d -dimensional j -reduction which, given a graph G' , forms G by deleting some vertex v of degree $j + d$ and then adding j edges between the vertices in $N_{G'}(v)$. Extensions and reductions of this type were first introduced by Henneberg [3], and so are also known as Henneberg moves and inverse Henneberg moves respectively.

Tay later used the inductive construction in Theorem 2.4 to obtain a new proof of Theorem 2.3, see [15]. Combining Theorems 2.3 and 2.4 gives the following result, which is the starting point of this paper:

Theorem 2.5. *A multi-graph $G = (V, E)$ is the edge-disjoint union of d spanning trees if and only if there exists a sequence of graphs*

$$K_1 = G^{(1)} \rightarrow G^{(2)} \rightarrow \dots \rightarrow G^{(n)} = G$$

such that for all $2 \leq i \leq n$, $G^{(i)}$ is obtained from $G^{(i-1)}$ by a d -dimensional j -extension, for some $0 \leq j \leq d - 1$.

3. INDUCTIVE CONSTRUCTION FOR d -TREE DECOMPOSITIONS

Given a d -tree decomposition $\mathcal{G} = (G; T_1, \dots, T_d)$, the construction in Theorem 2.5 can be adapted so that it not only constructs the multi-graph G , but simultaneously constructs each of the monochrome spanning trees T_1, T_2, \dots, T_d through a sequence of 1-dimensional Henneberg moves. This extension is not difficult to see, however to the best of our knowledge this result has not appeared in the literature. For completeness, we prove it here.

Lemma 3.1. *Let $T = (V, E)$ be a simple graph, and let $S = (U, F)$ be a subgraph of T which is a tree. Suppose $T' = (V', E')$ is formed from T by a 1-dimensional $|F|$ -extension which deletes the edges in F , and appends a vertex v incident to $|F| + 1$ new edges. Then T is a tree if and only if T' is a tree.*

Proof. By the definition of a d -dimensional j -extension, $|E| = |E'| - 1$ and $|V| = |V'| - 1$. It follows that $|E'| = |V'| - 1$ if and only if $|E| = |V| - 1$. So to prove that T is a tree if and only if T' is a tree, it remains to note that T is connected if and only if T' is connected. \square

The following elementary fact, which is a consequence of Theorem 2.3 and the definition of \mathcal{G}_d , shall also be useful:

Lemma 3.2. *Let G be a multi-graph, with $|V(G)| \geq 2$, which admits a d -tree decomposition for some $d \geq 1$. Then $d \leq \delta(G) \leq 2d - 1$.*

This lemma implies we only need to consider d -dimensional j -extensions of d -tree decompositions where $0 \leq j \leq d - 1$.

Definition 3.3. A d -tree decomposition $\mathcal{G}' = (G'; T'_1, \dots, T'_d)$ is said to be obtained from a d -tree decomposition $\mathcal{G} = (G; T_1, \dots, T_d)$ by a d -tree j -extension, for some $0 \leq j \leq d - 1$, if

- (a) G' is obtained from G by a d -dimensional j -extension, and,
- (b) for each $1 \leq i \leq d$, T'_i is obtained from T_i by a 1-dimensional k_i -extension, for some $0 \leq k_i \leq d - 1$.

Note that $\sum_{i=1}^d k_i = j$.

Proposition 3.4. *Let $\mathcal{G}' \in \mathcal{G}_d$ be a d -tree decomposition with $\mathcal{G}' \neq \mathcal{K}_1$. Then there exists a d -tree decomposition $\mathcal{G} \in \mathcal{G}_d$ such that \mathcal{G}' is a d -tree j -extension of \mathcal{G} .*

Proof. Let $\mathcal{G}' = (G'; T'_1, \dots, T'_d)$ where $G' = (V', E')$ and $T'_i = (V', E'_i)$ for each $i = 1, \dots, d$. By Lemma 3.2, there exists $v \in V'$ such that $d \leq d_{G'}(v) \leq 2d - 1$. For each tree T'_i , perform a Henneberg reduction at v which forms the graph $T_i = (V' - v, E_i)$ by deleting v and adding a set of edges F_i between the vertices in $N_{T'_i}(v)$ such that $(N_{T'_i}(v), F_i)$ is a tree.

Since T'_1, T'_2, \dots, T'_d are spanning trees of G' , we know $d_{T'_i}(v) \geq 1$ for all $1 \leq i \leq d$. By Lemma 3.2, we have that

$$\sum_{i=1}^d d_{T'_i}(v) = d_{G'}(v) \leq 2d - 1.$$

Hence $1 \leq d_{T'_i}(v) \leq d$ for all $1 \leq i \leq d$. Since $(N_{T'_i}(v), F_i)$ is a tree, we must have $|F_i| = |N_{T'_i}(v)| - 1 = d_{T'_i}(v) - 1$. In other words, the inverse Henneberg move which forms T_i from T'_i is a 1-dimensional k_i -reduction, with $k_i = d_{T'_i}(v) - 1$. Which implies $0 \leq k_i \leq d - 1$.

Let $G = (V' - v, E_1 \cup \dots \cup E_d)$. By Lemma 3.1, T_1, \dots, T_d are trees and so $\mathcal{G} = (G; T_1, \dots, T_d)$ is a d -tree decomposition. Finally, the move which forms G from G' deletes the vertex v , and then adds $\sum_{i=1}^d (d_{T'_i}(v) - 1) = d_{G'}(v) - d$ edges between the vertices in $N_{G'}(v)$. This is, by definition, a d -dimensional j -reduction, where $j = d_{G'}(v) - d$, and since $d \leq d_{G'}(v) \leq 2d - 1$, we have $0 \leq j \leq d - 1$. \square

Corollary 3.5. *Let $\mathcal{G} = (G; T_1, \dots, T_d)$ be a d -tree decomposition. Then, there exists a sequence of d -tree decompositions*

$$\mathcal{K}_1 = \mathcal{G}^{(1)} \rightarrow \mathcal{G}^{(2)} \rightarrow \dots \rightarrow \mathcal{G}^{(n)} = \mathcal{G}$$

such that for all $2 \leq i \leq n$, $\mathcal{G}^{(i)}$ is obtained from $\mathcal{G}^{(i-1)}$ by a d -tree j -extension, for some $0 \leq j \leq d - 1$.

4. REALISATIONS FOR 2-TREE DECOMPOSITIONS

In this section, we apply our inductive construction from Corollary 3.5 to show that every 2-tree decomposition has a plane realisation. To do this, we first extend the definitions of a well-positioned framework (G, p) in \mathbb{R}^d and an induced framework colouring κ_p , which were given in Section 2.2, to accommodate multi-graphs.

Let $G = (V, E)$ be a multi-graph with no loops and let $p : V \rightarrow \mathbb{R}^d$ be an injective map. As before we refer to the pair (G, p) as a *framework* in \mathbb{R}^d . Suppose a pair of vertices $u, v \in V$ are joined by exactly t edges, where $1 \leq t \leq d$. The pair $\{p(u), p(v)\}$ is said to be *well-positioned* if there exist exactly t distinct elements $j_1, \dots, j_t \in \{1, \dots, d\}$ such that

$$\|p(u) - p(v)\|_\infty = |(p(u) - p(v)) \cdot e_{j_k}|$$

for $k = 1, 2, \dots, t$. We refer to j_1, \dots, j_t as the *framework colours* for the pair $\{p(u), p(v)\}$. The framework (G, p) , on the multi-graph G , is said to be *well-positioned* if, for every edge $uv \in E(G)$, the pair $\{p(u), p(v)\}$ is well-positioned.

When a framework (G, p) is well-positioned, the multi-graph G inherits an edge-labelling $\kappa_p : E(G) \rightarrow \{1, 2, \dots, d\}$ whereby each of the t edges connecting a pair of vertices u and v is assigned one of the distinct framework colours j_1, \dots, j_t for the pair $\{p(u), p(v)\}$. The edge-labelling κ_p is referred to as a *framework colouring* for (G, p) . Note that this framework colouring is unique up to permutation of colours between parallel edges. Such permutations will not create any problems in what follows. The subgraph

$G[\kappa_p^{-1}(k)]$ spanned by edges with framework colour k is again referred to as a *monochrome subgraph* of G .

Given a d -tree decomposition $\mathcal{G} = (G; T_1, \dots, T_d)$ of a multi-graph G , a *realisation* of \mathcal{G} is a framework (G, p) in \mathbb{R}^d with the property that T_1, \dots, T_d are the induced monochrome subgraphs of G , for some choice of framework colouring κ_p . Note that by relabelling T_1, \dots, T_d we may assume, without loss of generality, that for $i = 1, \dots, d$, edges uv in T_i satisfy,

$$\|p(u) - p(v)\|_\infty = |(p(u) - p(v)) \cdot e_i|.$$

Proposition 4.1. *Let $\mathcal{G} = (G; T_1, T_2)$ be a 2-tree decomposition and suppose $\mathcal{G}' = (G'; T'_1, T'_2)$ is a 2-tree decomposition which is obtained by applying a 2-tree 0-extension to \mathcal{G} .*

If \mathcal{G} has a realisation p in the plane then \mathcal{G}' has a realisation p' in the plane with the property that $p'(w) = p(w)$ for all $w \in V(G)$.

Proof. Suppose the 2-dimensional 0-reduction which forms G from G' deletes the vertex v . By the definition of a 2-dimensional 0-reduction, $d_{G'}(v) = 2$, and G is formed from G' by deleting v and the two edges incident to v . Since T'_1 and T'_2 are both spanning trees of G' , this implies that v is a leaf node of both T'_1 and T'_2 . Hence $T_i = T'_i - v$ for $i \in \{1, 2\}$. Let $N_{G'}(v) = \{x, y\}$ where x and y may or may not be distinct. Without loss of generality, suppose $vx \in E(T'_1)$ and $vy \in E(T'_2)$. Since $G = G' - v$, we know that $(G' - v, p)$ is well-positioned. So let $p'(w) = p(w)$ for all $w \in V(G)$. For (G', p') to be well-positioned, it only remains to find a position for $p'(v)$ such that the pairs $\{p'(v), p'(x)\}$ and $\{p'(v), p'(y)\}$ are well-positioned.

If $x \neq y$ then $p'(x) \neq p'(y)$. If we place v within a sufficiently small distance of the intersection of the lines $p'(x) + \lambda(1, 0)$ and $p'(y) + \mu(0, 1)$, then $\{p'(v), p'(x)\}$ and $\{p'(v), p'(y)\}$ will have framework colours 1 and 2 respectively.

Now suppose $x = y$. Then v is met by a double-edge; one edge in T'_1 and the other in T'_2 . To satisfy the constraints induced by this colouring, we must place $p'(v)$ such that

$$\|p'(v) - p'(x)\|_\infty = |v_1 - x_1| = |v_2 - x_2|.$$

If we place $p'(v)$ at any point of the lines $p'(x) + \lambda(1, 1)$ or $p'(x) + \lambda(1, -1)$, then we will satisfy these constraints, and, since vx is a double-edge, the pair $\{p'(v), p'(x)\}$ will be well-positioned.

In both cases we can choose $p'(v)$ so that it is not coincident with another vertex of G' . Moreover, (G', p') is well-positioned, and the induced monochrome subgraphs are T'_1 and T'_2 . \square

To complete the inductive construction we require a geometric method of realising 2-tree 1-extensions which accommodates parallel edges.

Proposition 4.2. *Let $\mathcal{G} = (G; T_1, T_2)$ be a 2-tree decomposition and suppose $\mathcal{G}' = (G'; T'_1, T'_2)$ is a 2-tree decomposition which is obtained by applying a 2-tree 1-extension to \mathcal{G} .*

If \mathcal{G} , and every 2-tree decomposition with fewer vertices than \mathcal{G} , has a realisation p in the plane then \mathcal{G}' has a realisation p' in the plane.

Proof. Suppose the 2-dimensional 1-reduction which forms G from G' deletes the vertex v . By the definition of a 2-dimensional 1-reduction, $d_{G'}(v) = 3$, and G is formed from G' by deleting v and all three edges incident to v , before adding a single edge between the vertices in $N_{G'}(v)$. Since T'_1 and T'_2 are both spanning trees of G' , this implies that v is a leaf node of one of these trees, and is incident to two edges of the other tree. Without loss of generality assume $d_{T'_1}(v) = 1$ and $d_{T'_2}(v) = 2$. Let $N_{T'_1}(v) = \{x\}$ and $N_{T'_2}(v) = \{y, z\}$, where $y \notin \{x, z\}$, but potentially $x = z$. Then, $T_1 = T'_1 - v$. In order for T_2 to be a tree, it must be formed by deleting v from T'_2 , and then adding an edge e between y and z . In other words, $T_2 = T'_2 - v + e$. Let (G, p) be a realisation for \mathcal{G} in the plane. We shall use p to construct a realisation p' of G' .

For all pairs of points $(a, b), (c, d) \in \mathbb{R}^2$, write $(a, b) \leq (c, d)$ if $a \leq c$ and $b \leq d$ and write $(a, b) < (c, d)$ if, in addition, $(a, b) \neq (c, d)$. Note that reflecting the set of points $\{p(w) : w \in V(G)\}$ through either of the coordinate axes in \mathbb{R}^2 generates another realisation of \mathcal{G} in which the framework colours for all pairs $\{p(u), p(w)\}$ are preserved. Thus we may assume, without loss of generality, that $p(z) < p(y)$.

Case 1. Suppose $x \neq z$.

If $yz \notin E(T_1)$, then let $p'(w) = p(w)$ for all $w \in V(G)$. Since $yz \in E(T_2)$, any placement of $p'(v)$ on the line $p'(y) + \mu(p'(z) - p'(y))$, $\mu \in \mathbb{R}$, which is distinct from $p'(y)$ and $p'(z)$, will ensure the pairs $\{p'(v), p'(y)\}$ and $\{p'(v), p'(z)\}$ are well-positioned with framework colour 2. Similarly, for all $-1 < a < 1$, any placement of $p'(v)$ on the line $p'(x) + \lambda(1, a)$, $\lambda \in \mathbb{R}$, which is distinct from $p'(x)$, will ensure $\{p'(v), p'(x)\}$ is well-positioned with framework colour 1. Hence, if we place $p'(v)$ at the point of intersection of the two lines $p'(x) + \lambda(1, a)$ and $p'(y) + \mu(p'(z) - p'(y))$, and a is chosen so that this intersection is not coincident with any other vertex of G' , then (G', p') is a realisation of \mathcal{G}' . See Figure 2(a).

If $yz \in E(T_1)$, then proceed as follows: Let o denote the intersection of the lines $p(x) + \lambda(1, 0)$ and $p(y) + \mu(1, 1)$ and choose $\epsilon > 0$ and $\delta > 0$. For a pair of vertices $u, w \in V(G)$, write $u \sim w$ if either $u = w$, or, w is joined to u by a sequence of parallel edges in $G - e$.

If $o \leq p(z)$ then, informally, $p'(v)$ may be placed to the right of o and $p(z)$ translated upwards. See Figure 2(b). Any vertex in G which is joined to z by a sequence of parallel edges in $G - e$ must also be translated upwards. Formally, define p' by setting $p'(v) = o + \delta(1, 0)$ and, for each $w \in V(G)$,

$$p'(w) = \begin{cases} p(w) + \epsilon(0, 1) & \text{if } w \sim z, \text{ and,} \\ p(w) & \text{otherwise.} \end{cases}$$

If $p(z) < o < p(x)$ then, informally, place $p'(v)$ below $p(z)$ and then translate $p(z)$ upwards. See Figure 2(c). Any vertex in G which is joined to z by a sequence of parallel edges in $G - e$ must also be translated upwards. Formally, define p' by setting $p'(v) = p(z) + \delta(0, -1)$ and, for each $w \in V(G)$,

$$p'(w) = \begin{cases} p(w) + \epsilon(0, 1) & \text{if } w \sim z, \text{ and,} \\ p(w) & \text{otherwise.} \end{cases}$$




Figure 2. Some choices for $p'(v)$ in (G', p') when x, y and z are distinct. Features of (G', p') are shown in black, with solid and dashed black lines denoting edges in T_1' and T_2' respectively. Unless otherwise indicated, $p'(w) = p(w)$ for all $w \in V(G)$. For each vertex $w \in \{x, y, z\}$ we highlight in grey either a shaded region or a line, such that any placement P of v within this region ensures the pair $(P, p'(w))$ satisfies the geometric constraints implied by the graph colouring of G' . In each case, our chosen placement, $p'(v)$, lies within the intersection of these three regions. Note that in examples (b) and (c), both T_1 and T_2 contain a yz -edge, so to ensure the pair $(p'(y), p'(z))$ is well-positioned after $yz \in E(T_2)$ is deleted, we shift z upwards.

If $p(z) < o \leq p(y)$ and $p(x) < o$ then, informally, place $p'(v)$ above $p(y)$ and then translate $p(y)$ downwards. See Figure 3(a). Any vertex in G which is joined to y by a sequence of parallel edges in $G - e$ must also be translated downwards. Formally, define p' by setting $p'(v) = p(y) + \delta(0, 1)$ and, for each $w \in V(G)$,

$$p'(w) = \begin{cases} p(w) + \epsilon(0, -1) & \text{if } w \sim y, \text{ and,} \\ p(w) & \text{otherwise.} \end{cases}$$

If $p(z) < p(y) < o$ and $p(x) < o$ then, informally, place $p'(v)$ above o and then translate $p(y)$ downwards. See Figure 3(b). Any vertex in G which is joined to y by a sequence of parallel edges in $G - e$ must also be translated downwards. Formally, define p' by setting $p'(v) = o + \delta(0, 1)$ and, for each $w \in V(G)$,

$$p'(w) = \begin{cases} p(w) + \epsilon(0, -1) & \text{if } w \sim y, \text{ and,} \\ p(w) & \text{otherwise.} \end{cases}$$

If $p(z) < o$ and $p(x) = o$ then $p(x), p(y)$ and $p(z)$ are collinear. First suppose $p(z) < p(y) < p(x)$. Place $p'(v)$ at the intersection of the horizontal line through $p(x)$ and the vertical line through $p(y)$. Then translate $p(y)$ downwards. See Figure 3(c). Any vertex in G which is joined to y by a sequence of parallel edges in $G - e$ must also be translated downwards. Thus, for each $w \in V(G)$ we define,

$$p'(w) = \begin{cases} p(w) + \epsilon(0, -1) & \text{if } w \sim y, \text{ and,} \\ p(w) & \text{otherwise.} \end{cases}$$

Now suppose $p(z) < p(x) < p(y)$. In this case a convenient placement for v is not to hand. Instead, contract the parallel edges between y and




Figure 3. Further choices for $p'(v)$ in (G', p') when x, y and z are distinct and $yz \in E(T_1)$. In each case, we place $p'(y)$ below $p(y)$ to ensure the pair $(p'(y), p'(z))$ is well-positioned after $yz \in E(T_2)$ is deleted.

z to form a new 2-tree decomposition $\mathcal{G}^o = (G^o; T_1^o, T_2^o)$. Write $V(G^o) = V(G) - \{y, z\} + \{w_0\}$ where w_0 is the vertex obtained by identifying y and z . Since \mathcal{G}^o has fewer vertices than \mathcal{G} , there exists a realisation p^o for \mathcal{G}^o in the plane. From this realisation, we can construct a new realisation p for \mathcal{G} with the property that $p(z) < p(y) < p(x)$. Formally, set $p(z) = p^o(w_0)$, $p(y) = p^o(w_0) + \epsilon'(1, 1)$ for some sufficiently small $\epsilon' > 0$, and $p(w) = p^o(w)$ for all remaining $w \in V(G)$. Now proceed as above.

Case 2. Suppose $x = z$.

In this case v sends a double edge to x : one edge in T_1' and the other in T_2' . If $xy \notin E(T_1)$, then let $p'(w) = p(w)$ for all $w \in V(G)$. Place $p'(v)$ at the intersection of the lines $p'(x) + \lambda(1, 1)$ and $p'(y) + \mu(a, 1)$ where $-1 < a < 1$. Choose a such that $p'(v)$ is not coincident with any other vertex of (G', p') . See Figure 4(a).

If $xy \in E(T_1)$ then, informally, place $p'(v)$ at $p(y)$ and then translate $p(y)$ downwards. Any vertex in G which is joined to y by a sequence of parallel edges in $G - e$ must also be translated downwards. Formally, define p' by setting $p'(v) = p(y)$ and, for each $w \in V(G)$,

$$p'(w) = \begin{cases} p(w) + \epsilon(0, -1) & \text{if } w \sim y, \text{ and,} \\ p(w) & \text{otherwise,} \end{cases}$$

where $\epsilon > 0$. See Figure 4(b)

In each case, ϵ and δ can be chosen sufficiently small so that $p'(v)$ satisfies the required constraints and so that the points $\{p'(w) : w \in V(G')\}$ are distinct. Thus (G', p') is a realisation for \mathcal{G}' . \square

We can now prove that every 2-tree decomposition has a realisation in the plane.

Theorem 4.3. *Let $\mathcal{G} = (G; T_1, T_2)$ be a 2-tree decomposition. Then there exists a realisation for \mathcal{G} in the plane.*

Proof. Corollary 3.5 gives a sequence of 2-tree decompositions and 2-tree 0 and 1-extensions which construct \mathcal{G} from the base element \mathcal{K}_1 . Note that




Figure 4. Choices for $p'(v)$ in (G', p') when $x = z$. Since both T'_1 and T'_2 contain an xv -edge, the line l_{xv} through $p'(x)$ and $p'(v)$ has slope 1.

a realisation of \mathcal{K}_1 is obtained by placing the vertex of \mathcal{K}_1 anywhere in the plane. By Propositions 4.1 and 4.2, there exists a realisation for every 2-tree decomposition in this sequence, in particular, such a realisation exists for \mathcal{G} . \square

As a corollary we obtain an alternative proof of the following result from [8].


Corollary 4.4. *Let G be a $(2, 2)$ -tight simple graph. Then there exists a placement p such that (G, p) is well-positioned and minimally rigid in $(\mathbb{R}^2, \|\cdot\|_\infty)$.*

Proof. By Theorem 2.3, G admits a 2-tree decomposition $\mathcal{G} = (G; T_1, T_2)$. By Theorem 4.3, this 2-tree decomposition has a realisation (G, p) in the plane. By Theorem 2.2, this realisation is minimally rigid in $(\mathbb{R}^2, \|\cdot\|_\infty)$. \square


5. SYMMETRIC 2-TREE DECOMPOSITIONS

In this section we adapt the methods of the previous pages to show that every *symmetric* 2-tree decomposition, with no fixed edges, can be realised as a symmetric framework in the plane which is minimally rigid with respect to ℓ^∞ distance constraints. We focus on frameworks with reflectional symmetry through a coordinate axis. Our motivation comes from recent work ([9]) which characterises the class of symmetric graphs which admit a symmetric and minimally rigid realisation in the plane, in terms of certain symmetric 2-tree decompositions. The main results of this section make the task of constructing examples of symmetric minimally rigid frameworks significantly easier. These results also suggest that further adaptations may be possible in other contexts, such as gain graph constructions for symmetric frameworks.

By a \mathbb{Z}_2 -*symmetric multi-graph* we will mean a pair (G, θ) consisting of a multi-graph G , with no loops, and a non-trivial group homomorphism $\theta : \mathbb{Z}_2 \rightarrow \text{Aut}(G)$. Let $\mathbb{Z}_2 = \langle s \rangle$. To simplify notation, we denote $\theta(s)$ by s_θ and for each edge $e = v_1v_2$ we write $s_\theta(e) = s_\theta(v_1)s_\theta(v_2)$. The *vertex orbit* of $v \in V(G)$ is the pair of vertices $\{v, s_\theta(v)\}$ and the *edge orbit* of $e \in E(G)$ is $\{e, s_\theta(e)\}$. We say that a vertex v (respectively an edge e) is *fixed* if $v = s_\theta(v)$ (respectively $e = s_\theta(e)$).



(a) A symmetric 2-tree decomposition \mathcal{G}' formed from a symmetric 2-tree decomposition \mathcal{G} by a symmetric 2-tree 0-extension.



(b) Two symmetric 2-tree decompositions, \mathcal{G}' and \mathcal{G}'' , each formed by a symmetric 2-tree 0-extension on the symmetric 2-tree decomposition \mathcal{G} .

Figure 5. Symmetric 2-tree 0-extensions under half-turn (a) and single mirror (b) symmetry.

Definition 5.1. A *symmetric 2-tree decomposition* is a tuple $\mathcal{G} = (G; T_1, T_2; \theta)$ such that,

- (a) $(G; T_1, T_2)$ is a 2-tree decomposition,
- (b) (G, θ) is a \mathbb{Z}_2 -symmetric multi-graph, and
- (c) $s_\theta(T_1) = T_1$ and $s_\theta(T_2) = T_2$.

Denote by \mathcal{G}_2^{sym} the set of all symmetric 2-tree decompositions with no fixed edges. We formally include in \mathcal{G}_2^{sym} the tuple $\mathcal{K}_1 = (K_1; T_1, T_2; \theta)$ where K_1 is the graph with a single vertex v_0 and no edges, T_1 and T_2 have empty edge set, and $\theta : \mathbb{Z}_2 \rightarrow \text{Aut}(K_1)$ is the trivial group homomorphism with $s_\theta(v_0) = v_0$. This tuple will form the base element of a construction scheme for \mathcal{G}_2^{sym} .

Lemma 5.2. Let $\mathcal{G} = (G; T_1, T_2; \theta) \in \mathcal{G}_2^{sym}$ with $\mathcal{G} \neq \mathcal{K}_1$. Then s_θ fixes exactly one vertex which has even degree at least 4.

Proof. The case where G is a simple graph is proved in [9, Lemma 3] and this proof extends to the multi-graph case. \square

5.1. Multi-graph construction scheme for \mathcal{G}_2^{sym} . The following two graph moves were applied in the context of simple graphs, and with $d = 2$, in [9] (where they were referred to as \mathbb{Z}_2 -symmetric 1-extensions and \mathbb{Z}_2 -symmetric 2-extensions). Here we will first extend the moves to multi-graphs, allowing $d \in \{1, 2\}$, and then introduce corresponding moves for symmetric 2-tree decompositions.

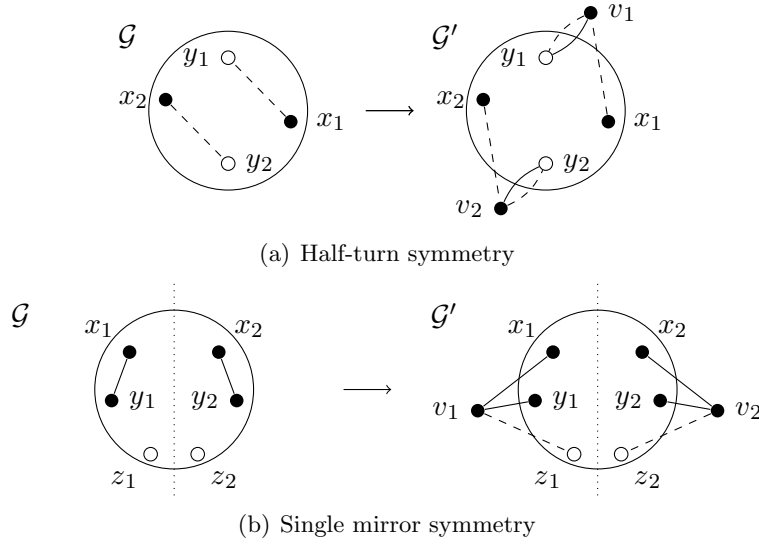


Figure 6. Symmetric 2-tree 1-extensions. In each case, the symmetric 2-tree decomposition \mathcal{G}' is obtained from the symmetric 2-tree decomposition \mathcal{G} by a symmetric 2-tree 1-extension.

Definition 5.3. A \mathbb{Z}_2 -symmetric multi-graph (G', θ') is said to be obtained from a \mathbb{Z}_2 -symmetric multi-graph (G, θ) by a *symmetric d -dimensional 0-extension* if,

- (a) $V(G') = V(G) \cup \{v, s_{\theta'}(v)\}$ where $v, s_{\theta'}(v) \notin V(G)$ and $v \neq s_{\theta'}(v)$,
- (b) $s_{\theta'}|_{V(G)} = s_\theta$,
- (c) $E(G') = E(G) + \{vv_i, s_{\theta'}(vv_i) : i = 1, \dots, d\}$ for some $v_1, \dots, v_d \in V(G)$ not necessarily distinct.

See Figure 5 for examples of such a move when $d = 2$.

Definition 5.4. A \mathbb{Z}_2 -symmetric multi-graph (G', θ') is said to be obtained from a \mathbb{Z}_2 -symmetric multi-graph (G, θ) by a *symmetric d -dimensional 1-extension* if,

- (a) $V(G') = V(G) \cup \{v, s_{\theta'}(v)\}$ where $v, s_{\theta'}(v) \notin V(G)$ and $v \neq s_{\theta'}(v)$,
- (b) $s_{\theta'}|_{V(G)} = s_\theta$,
- (c) there exist $d + 1$ vertices $v_1, \dots, v_{d+1} \in V(G)$, with $e = v_1v_2 \in E(G)$ but which are otherwise not necessarily distinct, such that $E(G') = E(G) - \{e, s_\theta(e)\} + \{vv_i, s_{\theta'}(vv_i) : i = 1, \dots, d + 1\}$.

See Figure 6 for examples when $d = 2$.

We now adapt the above symmetric graph moves to incorporate symmetric 2-tree decompositions. Again see Figures 5 and 6 for illustrations of these moves.

Definition 5.5. A symmetric 2-tree decomposition $\mathcal{G}' = (G'; T'_1, T'_2; \theta')$ is said to be obtained from a symmetric 2-tree decomposition $\mathcal{G} = (G; T_1, T_2; \theta)$ by a *symmetric 2-tree j -extension*, where $j \in \{0, 1\}$, if

- (a) (G', θ') is obtained from (G, θ) by a symmetric 2-dimensional j -extension, and,

- (b) for $i \in \{1, 2\}$, (T'_i, θ') is obtained from (T_i, θ) by a symmetric 1-dimensional k_i -extension, for some $k_i \in \{0, 1\}$ where $k_1 + k_2 = j$.

Reconsider Figures 5 and 6, this time noting the edge colourings.

We now prove the existence of a construction scheme for symmetric 2-tree decompositions in \mathcal{G}_2^{sym} which uses symmetric 2-tree 0 and 1-extensions.

Theorem 5.6. *Let $\mathcal{G}' = (G'; T'_1, T'_2; \theta') \in \mathcal{G}_2^{sym}$ with $\mathcal{G}' \neq \mathcal{K}_1$. Then there exists $\mathcal{G} = (G; T_1, T_2; \theta) \in \mathcal{G}_2^{sym}$ such that \mathcal{G}' is obtained from \mathcal{G} by either a symmetric 2-tree 0-extension or a symmetric 2-tree 1-extension.*

Proof. By Lemma 3.2, G' has a vertex v of degree 2 or 3, and by Lemma 5.2, this vertex is not fixed. Hence the vertex orbit $\{v, s_{\theta'}(v)\}$ contains two distinct vertices with $d_{G'}(v) = d_{G'}(s_{\theta'}(v))$. Since G' is the edge-disjoint union of spanning trees T'_1 and T'_2 , v is incident to at least one edge $e_i \in E(T'_i)$ from each of these trees. Since $d_{G'}(v) \leq 3$, either v is incident to one other edge, e_3 , or v is incident to no other edges.

Claim 5.7. If v is incident to no other edges, then \mathcal{G}' is formed from a symmetric 2-tree decomposition \mathcal{G} by a symmetric 2-tree 0-extension.

Proof. In this case $N_{G'}(v) = \{u_1, u_2\}$, where u_1 and u_2 need not be distinct. Since $d_{G'}(v) = 2$, and $s_{\theta'}(v) \notin N_{G'}(v)$, we can remove both v and $s_{\theta'}(v)$ from G' by separate 2-dimensional 0-reductions to obtain a \mathbb{Z}_2 -symmetric multi-graph (G, θ) with $s_\theta = s_{\theta'}|_{V(G)}$. Let $T_i = T'_i - \{v, s_{\theta'}(v)\}$ and note that $s_\theta(T_i) = T_i$. Thus $\mathcal{G} = (G; T_1, T_2; \theta) \in \mathcal{G}_2^{sym}$. Further \mathcal{G}' may be obtained from \mathcal{G} by a symmetric 2-tree 0-extension. \square

Claim 5.8. If v is incident to a third edge e_3 , then \mathcal{G}' is formed from a symmetric 2-tree decomposition \mathcal{G} by a symmetric 2-tree 1-extension.

Proof. Without loss of generality, suppose $e_3 \in E(T'_1)$. For $i \in \{1, 2, 3\}$, none of the edges e_i are fixed, so they each terminate at some $u_i \in V(G') - \{v, s_{\theta'}(v)\}$. Hence each edge $s_{\theta'}(e_i)$ incident to $s_{\theta'}(v)$ terminates at $s_{\theta'}(u_i) \in V(G') - \{v, s_{\theta'}(v)\}$.

Since $e_1, e_3 \in E(T'_1)$ and T'_1 is a tree, $u_1 u_3 \notin E(T'_1)$, and so, since T'_1 is symmetric under θ' , $s_{\theta'}(u_1) s_{\theta'}(u_3) \notin E(T'_1)$ either. If $\{u_1, u_3\} = \{s_{\theta'}(u_1), s_{\theta'}(u_3)\}$ then $v, u_1, s_{\theta'}(v), u_3, v$ is a cycle in T'_1 , which is a contradiction. Hence $\{u_1, u_3\} \neq \{s_{\theta'}(u_1), s_{\theta'}(u_3)\}$ and so we can perform a 2-dimensional 1-reduction at v which adds the edge $u_1 u_3$, followed by a 2-dimensional 1-reduction at $s_{\theta'}(v)$ which adds the edge $s_{\theta'}(u_1) s_{\theta'}(u_3)$ to form the graph G .

Let $\mathcal{G} = (G; T_1, T_2; \theta)$ be the symmetric 2-tree decomposition with $G = G' - \{v, s_{\theta'}(v)\} + \{u_1 u_3, s_{\theta'}(u_1 u_3)\}$, $T_1 = T'_1 - \{v, s_{\theta'}(v)\} + \{u_1 u_3, s_{\theta'}(u_1 u_3)\}$, $T_2 = T'_2 - \{v, s_{\theta'}(v)\}$ and $s_\theta = s_{\theta'}|_{V(G)}$. Then $s_\theta(T_i) = T_i$ and so $\mathcal{G} \in \mathcal{G}_2^{sym}$. Further T_1 was obtained from T'_1 by a pair of 1-dimensional 1-reductions, and T_2 was formed from T'_2 by a pair of 1-dimensional 0-reductions. Thus we can reconstruct \mathcal{G}' from \mathcal{G} by a symmetric 2-tree 1-extension. \square

\square

This result implies the inductive constructions sought:

Corollary 5.9. *Let $\mathcal{G} = (G; T_1, T_2; \theta) \in \mathcal{G}_2^{sym}$. Then, there exists a sequence of symmetric 2-tree decompositions in \mathcal{G}_2^{sym} ,*

$$\mathcal{K}_1 = \mathcal{G}^{(1)} \rightarrow \mathcal{G}^{(2)} \rightarrow \dots \rightarrow \mathcal{G}^{(n)} = \mathcal{G}$$

such that for all $2 \leq i \leq n$, $\mathcal{G}^{(i)}$ is obtained from $\mathcal{G}^{(i-1)}$ by a symmetric 2-tree j -extension, for some $j \in \{0, 1\}$.

5.2. Realisations with \mathcal{C}_s -symmetry. We shall now show how to construct a symmetric realisation for any symmetric 2-tree decomposition from the class \mathcal{G}_2^{sym} . We prove this explicitly for realisations under reflection symmetry. The argument for half-turn symmetry is similar.

can construct a realisation with either reflectional symmetry or rotational symmetry.

A *symmetric placement* of a \mathbb{Z}_2 -symmetric multi-graph (G, θ) in the plane is a pair (p, τ) consisting of an injective map $p : V(G) \rightarrow \mathbb{R}^2$ and a representation $\tau : \mathbb{Z}_2 \rightarrow \text{GL}(\mathbb{R}^2)$ such that $\tau(s)(p(v)) = p(s_\theta(v))$ for all $v \in V(G)$. If $\tau(s)$ is a reflection in a coordinate axis then we refer to the pair (p, τ) as a \mathcal{C}_s -placement of (G, θ) . If $\tau(s)$ is a half-turn rotation about the origin then we refer to (p, τ) as a \mathcal{C}_2 -placement of (G, θ) .

A \mathcal{C}_s -realisation (respectively, \mathcal{C}_2 -realisation) for a symmetric 2-tree decomposition $\mathcal{G} = (G; T_1, T_2; \theta)$ in the plane is a \mathcal{C}_s -placement (respectively, \mathcal{C}_2 -placement) (p, τ) of (G, θ) with the property that (G, p) is a realisation for the 2-tree decomposition $(G; T_1, T_2)$.

Proposition 5.10. *Let $\mathcal{G} = (G; T_1, T_2; \theta)$ and $\mathcal{G}' = (G'; T'_1, T'_2; \theta')$ be a pair of symmetric 2-tree decompositions and suppose \mathcal{G}' is obtained by applying a symmetric 2-tree 0-extension to \mathcal{G} .*

If \mathcal{G} has a \mathcal{C}_s -realisation (p, τ) in the plane then \mathcal{G}' has a \mathcal{C}_s -realisation (p', τ) in the plane with the property that $p'(w) = p(w)$ for all $w \in V(G)$.

Proof. Let (p, τ) be a \mathcal{C}_s -realisation for \mathcal{G} in the plane. Suppose the symmetric 2-tree 0-extension which forms \mathcal{G}' from \mathcal{G} adjoins the vertices v and $s_{\theta'}(v)$ to G . Let \mathcal{G}'' be the intermediate (and non-symmetric) 2-tree decomposition obtained by deleting $s_{\theta'}(v)$ and its incident edges from \mathcal{G}' . Note that \mathcal{G}'' is obtained from \mathcal{G} by a (non-symmetric) 2-tree 0-extension. By Proposition 4.1, there exists a realisation p'' for \mathcal{G}'' with the property that $p''(w) = p(w)$ for all $w \in V(G)$. Define $p'(w) = p(w)$ for all $w \in V(G)$, $p'(v) = p''(v)$ and $p'(s_{\theta'}(v)) = \tau(s)(p'(v))$. The resulting pair (p', τ) is a \mathcal{C}_s -realisation for \mathcal{G}' . \square

To complete the inductive construction we now consider geometric placements for symmetric 2-tree 1-extensions.

Proposition 5.11. *Let $\mathcal{G} = (G; T_1, T_2; \theta)$ and $\mathcal{G}' = (G'; T'_1, T'_2; \theta')$ be a pair of symmetric 2-tree decompositions and suppose \mathcal{G}' is obtained by applying a symmetric 2-tree 1-extension to \mathcal{G} .*

If \mathcal{G} , and every symmetric 2-tree decomposition with fewer vertices than \mathcal{G} , has a \mathcal{C}_s -realisation in the plane then \mathcal{G}' has a \mathcal{C}_s -realisation in the plane.

Proof. Let (p, τ) be a \mathcal{C}_s -realisation for \mathcal{G} in the plane. We may assume, without loss of generality, that $\tau(s)$ is a reflection in the y -axis.

Suppose the symmetric 2-tree 1-extension which forms G' from G adjoins the vertices v and $s_{\theta'}(v)$. Then $d_{G'}(v) = 3$, and G is formed from G' by deleting v , $s_{\theta'}(v)$, and all edges incident to either of these vertices, before adding an edge e between the vertices in $N_{G'}(v)$ and another edge $s_{\theta'}(e)$ between the vertices in $N_{G'}(s_{\theta'}(v))$.

Suppose $d_{T_1'}(v) = 1$ and $d_{T_2'}(v) = 2$. For a pair of vertices $u, w \in V(G)$, write $u \sim^s w$ if either $u \in \{w, s_{\theta}(w)\}$, or, u is joined to either w or $s_{\theta}(w)$ by a sequence of parallel edges in $G - \{e, s_{\theta}(e)\}$. The construction of p' now follows the proof of Proposition 4.2 almost verbatim by replacing \sim with \sim^s and setting $p'(s_{\theta'}(v)) = \tau(s)(p'(v))$. The case where $d_{T_1'}(v) = 2$ and $d_{T_2'}(v) = 1$ can be proved by similar methods. \square

Theorem 5.12. *Let $\mathcal{G} = (G; T_1, T_2; \theta) \in \mathcal{G}_2^{sym}$ be a symmetric 2-tree decomposition with $\mathcal{G} \neq \mathcal{K}_1$. Then \mathcal{G} has a \mathcal{C}_s -realisation in the plane.*

Proof. Use Propositions 5.10 and 5.11, and apply a similar argument to the proof of Theorem 4.3. \square

Theorem 5.12 shows that it is always possible to construct examples of isostatic \mathcal{C}_s -symmetric frameworks in the ℓ^∞ plane which induce prescribed symmetric monochrome spanning trees with no fixed edges. In the following, note that not all $(2, 2)$ -tight \mathbb{Z}_2 -symmetric graphs admit a symmetric 2-tree decomposition.

Corollary 5.13. *Let (G, θ) be a \mathbb{Z}_2 -symmetric simple graph. If (G, θ) admits a symmetric 2-tree decomposition $\mathcal{G} = (G; T_1, T_2; \theta)$, with no fixed edges, then there exists a \mathcal{C}_s -realisation for \mathcal{G} in the plane which is well-positioned and minimally rigid in $(\mathbb{R}^2, \|\cdot\|_\infty)$.*

Proof. By Theorem 5.12, the symmetric 2-tree decomposition \mathcal{G} has a \mathcal{C}_s -realisation in the plane. By Theorem 2.2, this realisation is minimally rigid in $(\mathbb{R}^2, \|\cdot\|_\infty)$. \square

6. OPEN PROBLEMS

In Section 4, we showed that given any multi-graph with a partition of its edge set into two spanning trees, there exists a realisation of this 2-tree decomposition in the plane. It is not known whether this result extends to d -dimensions. Indeed, a solution here would settle a particular case of another open problem which is to determine whether every (d, d) -tight graph has a rigid placement in $(\mathbb{R}^d, \|\cdot\|_q)$ for $d \geq 3$ and $q \neq 2$.

Open Problem 6.1 (Rigidity for ℓ^q norms.). *Let G be a simple graph which is an edge-disjoint union of d spanning trees T_1, \dots, T_d , where $d \geq 3$.*

- (a) *Does there exist a placement of G in \mathbb{R}^d such that the induced monochrome subgraphs of G are precisely T_1, \dots, T_d ?*
- (b) *Does there exist an isostatic placement of G in $(\mathbb{R}^d, \|\cdot\|_q)$ for all (or for some) $q \neq 2$?*

A positive answer to (a) would imply a positive answer to (b) in the case $q = \infty$. If $d = 2$, then the answer to both questions is “yes”.

The following example suggests a different approach is required to extend Corollary 4.4 to higher dimensions.

Example 6.2. Suppose $\mathcal{G} = (G; T_1, T_2, T_3)$ is a 3-tree decomposition with a realisation (G, p) in \mathbb{R}^3 . Suppose x, y, z are vertices of G with $p(x) = (0, 0, 0)$, $p(y) = (-1, 3, -1)$ and $p(z) = (-1, 10, -3)$. Now suppose a 3-tree 0-extension is applied to \mathcal{G} at the vertices x, y, z , which adds the vertex v and results in a 3-tree decomposition $\mathcal{G}' = (G'; T'_1, T'_2, T'_3)$ with $xv \in T'_1$, $yv \in T'_2$ and $zv \in T'_3$. For these trees to correspond to the monochromatic trees induced by a framework colouring for G' , we must have

$$\begin{aligned} \|p(v) - p(x)\|_\infty &= |v_1 - x_1| = |v_1|, \\ \|p(v) - p(y)\|_\infty &= |v_2 - y_2| = |v_2 - 3|, \text{ and} \\ \|p(v) - p(z)\|_\infty &= |v_3 - z_3| = |v_3 + 3|. \end{aligned}$$

However, there is no such point $p(v) = (v_1, v_2, v_3)$ in \mathbb{R}^3 .

A related problem is that of constructing examples of redundantly rigid frameworks in $(\mathbb{R}^d, \|\cdot\|_q)$. Here a framework is *redundantly rigid* if it is rigid and every subframework obtained by the removal of a single edge is also rigid. Such frameworks have played a key role in the study of *global rigidity* for frameworks in Euclidean space (see for example [4]).

Open Problem 6.3 (Redundant rigidity for ℓ^q norms.). *Let G be a simple graph which is an edge-disjoint union of d Hamilton cycles H_1, \dots, H_d , where $d \geq 2$.*

- (a) *Does there exist a placement of G in \mathbb{R}^d such that the induced monochrome subgraphs of G are precisely H_1, \dots, H_d ?*
- (b) *Does there exist a redundantly rigid placement of G in $(\mathbb{R}^d, \|\cdot\|_q)$ for all (or for some) $q \neq 2$?*

A positive answer to (a) would imply a positive answer to (b) in the case $q = \infty$.

Similar realisation problems arise for other norms. For example, it is shown in [7] that rigidity for the cylinder norm on \mathbb{R}^3 is characterised by an induced framework colouring which decomposes the graph into an edge-disjoint union of a spanning tree and a spanning Laman graph. Again whether the existence of such a decomposition implies the existence of a geometric realisation is open.

Open Problem 6.4 (Rigidity for the cylinder norm.). *Let G be a simple graph which is an edge-disjoint union of two spanning subgraphs T and L where T is a tree and L is a Laman graph.*

- (a) *Does there exist a placement of G in $(\mathbb{R}^3, \|\cdot\|_{cyl})$ such that the induced monochrome subgraphs of G are precisely T and L ?*
- (b) *Does there exist an isostatic placement of G in $(\mathbb{R}^3, \|\cdot\|_{cyl})$?*

A positive answer to (a) would imply a positive answer to (b). The smallest graph in this class is $K_6 - e$, obtained by removing a single edge from the complete graph K_6 , and this graph does admit an isostatic placement in $(\mathbb{R}^3, \|\cdot\|_{cyl})$ (see [7]).

Realisation problems of this type also arise in considering *forced symmetric rigidity* (see for example [5, 11, 13] for the Euclidean context). A characterisation is obtained in [10] for forced reflectional symmetry in the

ℓ^∞ plane which is expressed in terms of framework colourings on the associated gain graphs. In this case, the gain graph is expressed as an edge-disjoint union of a spanning unbalanced map graph and a spanning tree.

Open Problem 6.5 (Forced symmetric rigidity for the ℓ^∞ norm.). *Let G_0 be a gain graph for a \mathbb{Z}_2 -symmetric graph which is expressible as an edge-disjoint union of two spanning subgraphs T and M , where T is a tree and M is an unbalanced map graph.*

Does there exist a placement of the covering graph G in the plane, with reflectional symmetry, such that the induced monochrome subgraphs of G_0 are precisely T and M ?

REFERENCES

- [1] András Frank and László Szegő. Constructive characterizations for packing and covering with trees. *Discrete Appl. Math.*, 131(2):347–371, 2003. Submodularity.
- [2] J. Graver, B. Servatius, and H. Servatius. *Combinatorial Rigidity*. Graduate Studies in Mathematics, 2. American Mathematical Society, Providence, RI., 1993.
- [3] L. E. L. Henneberg. *Die graphische Statik der starren Systeme*. B.G. Teubners Sammlung von Lehrbüchern auf dem Gebiete der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen. B. G. Teubner, 1911.
- [4] Bill Jackson and Tibor Jordán. Connected rigidity matroids and unique realizations of graphs. *Journal of Combinatorial Theory, Series B*, 94(1):1 – 29, 2005.
- [5] Tibor Jordán, Viktória E. Kaszanitzky, and Shin-ichi Tanigawa. Gain-sparsity and symmetry-forced rigidity in the plane. *Discrete & Computational Geometry*, 55(2):314–372, 2016.
- [6] D. Kitson. Finite and infinitesimal rigidity with polyhedral norms. *Discrete Comput. Geom.*, 54(2):390–411, 2015.
- [7] D. Kitson and R. L. Levene. Graph rigidity for unitarily invariant matrix norms. arxiv.org/abs/1709.08967, 2017.
- [8] D. Kitson and S. C. Power. Infinitesimal rigidity for non-Euclidean bar-joint frameworks. *Bull. Lond. Math. Soc.*, 46(4):685–697, 2014.
- [9] D. Kitson and B. Schulze. Symmetric isostatic frameworks with ℓ^1 or ℓ^∞ distance constraints. *Electron. J. Combin.*, 23(4), 2016.
- [10] D. Kitson and B. Schulze. Motions of grid-like reflection frameworks. *Journal of Symbolic Computation*, To appear.
- [11] Justin Malestein and Louis Theran. Frameworks with forced symmetry i: Reflections and rotations. *Discrete & Computational Geometry*, 54(2):339–367, 2015.
- [12] C. St. J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36:445–450, 1961.
- [13] Bernd Schulze and Shin ichi Tanigawa. Infinitesimal rigidity of symmetric bar-joint frameworks. *SIAM J. Discrete Math.*, 29:1259–1286, 2015.
- [14] T.-S. Tay. Rigidity of multigraphs. I. Linking rigid bodies in n -space. *J. Combin. Theory Ser. B*, 36(1):95–112, 1984.
- [15] T.-S. Tay. Rigidity of multigraphs. II. In *Graph theory, Singapore 1983*, volume 1073 of *Lecture Notes in Math.*, pages 129–134. Springer, Berlin, 1984.
- [16] W. T. Tutte. On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.*, 36:221–230, 1961.