Egerváry Research Group
on Combinatorial Optimization

Technical reportS

TR-2019-01. Published by the Egerváry Research Group, Pázmány P. sétány 1/C, H–1117, Budapest, Hungary. Web site: www.cs.elte.hu/egres. ISSN 1587–4451.

Minimum Cost Globally Rigid Subgraphs

Tibor Jordán and András Mihálykó

January 2019

Minimum Cost Globally Rigid Subgraphs

Tibor Jordán⋆ and András Mihálykó⋆⋆

Abstract

A d-dimensional framework is a pair (G, p), where $G = (V, E)$ is a graph and p is a map from V to $\mathbb{R}^d$. The length of an edge of G is equal to the distance between the points corresponding to its end-vertices. The framework is said to be globally rigid if its edge lengths uniquely determine all pairwise distances in the framework. A graph G is called globally rigid in $\mathbb{R}^d$ if every generic d-dimensional framework (G, p) is globally rigid. Global rigidity has applications in wireless sensor network localization, molecular conformation, formation control, CAD, and elsewhere.
Motivated by these applications we consider the following optimization problem: given a graph $G = (V, E)$, a non-negative cost function $c : E \longrightarrow \mathbb{R}_+$ on the edge set of G, and a positive integer d. Find a subgraph $H = (V, E')$ of G, on the same vertex set, which is globally rigid in $\mathbb{R}^d$ and for which the total cost $c(E') := \sum_{e \in E'} c(e)$ of the edges is as small as possible. This problem is NP-hard for all $d \geq 1$, even if c is uniform or G is complete and c is metric. We focus on the two-dimensional case, where we give $\frac{3}{2}$-approximation (resp. 2-approximation) algorithms for the uniform cost and metric versions. We also develop a constant factor approximation algorithm for the metric version of the d-dimensional problem, for every $d \geq 3$.

1 Introduction

A d-dimensional *framework* is a pair (G, p), where $G = (V, E)$ is a graph and p is a map from V to $\mathbb{R}^d$. We also call (G, p) a *realization* of G in $\mathbb{R}^d$. Two realizations (G, p) and (G, q) are *equivalent* if $||p(u) - p(v)|| = ||q(u) - q(v)||$ holds for all pairs u, v with $uv \in E$, where $||.||$ denotes the Euclidean norm in $\mathbb{R}^d$. The frameworks (G, p) and (G, q) are *congruent* if $||p(u) - p(v)|| = ||q(u) - q(v)||$ holds for all pairs u, v with $u, v \in V$. This is the same as saying that (G, q) can be obtained from (G, p) by an isometry of $\mathbb{R}^d$.

We say that (G, p) is *globally rigid* in $\mathbb{R}^d$ if every d-dimensional realization (G, q) of G which is equivalent to (G, p), is congruent to (G, p). In other words, the framework is globally rigid if its edge lengths uniquely determine all pairwise distances. This property makes the notion of global rigidity a fundamental concept in problems where we are given

⋆Department of Operations Research, Eötvös University, and the MTA-ELTE Egerváry Research Group on Combinatorial Optimization, Pázmány Péter sétány 1/C, 1117 Budapest, Hungary. Email: `jordan@cs.elte.hu`

⋆⋆Department of Operations Research, Eötvös University, Pázmány Péter sétány 1/C, 1117 Budapest, Hungary. Email: `mihalyko@cs.elte.hu`

partial information on the pairwise distances between pairs of a finite point set and our goal is to determine the configuration of the points, up to trivial transformations, see Section 1.3 below.

Saxe [27] showed that it is NP-hard to decide if even a 1-dimensional framework is globally rigid. The analysis and characterization of globally rigid frameworks become more tractable if we consider *generic frameworks*, i.e. frameworks (G, p) for which the set of coordinates of the points $p(v)$, $v \in V(G)$, is algebraically independent over the rationals. Results of Connelly [6] and Gortler, Healy and Thurston [12] imply that the global rigidity of a generic framework (G, p) in $\mathbb{R}^d$ depends only on the graph G, for all $d \geq 1$. Hence we may define a graph G to be *globally rigid* in $\mathbb{R}^d$ if every (or equivalently, if some) generic realization of G in $\mathbb{R}^d$ is globally rigid. The problem of finding a polynomially verifiable characterization for graphs which are globally rigid in $\mathbb{R}^d$ has been solved for $d = 1, 2$, but is a major open problem when $d \geq 3$.

1.1 The minimum cost globally rigid subgraph problem

In this paper we consider the following algorithmic problem. The input is a graph $G = (V, E)$, a non-negative cost function $c : E \longrightarrow \mathbb{R}_+$ on the edge set of G, and a positive integer d. The task is to find a subgraph $H = (V, E')$ of G, on the same vertex set, which is globally rigid in $\mathbb{R}^d$ and for which the total cost $c(E') := \sum_{e \in E'} c(e)$ of the edges is as small as possible.

We call this optimization problem the *Minimum cost globally rigid spanning subgraph problem* (or MCGRSS, for short). We shall focus on the following special cases of this problem: (i) if c is uniform, the goal is to find a *minimum size* globally rigid spanning subgraph (ii) in the *metric* MCGRSS problem the input graph G is complete and c satisfies the triangle inequality.

The Minimum cost globally rigid spanning subgraph problem (already in the special cases mentioned above) is NP-hard for all $d \geq 1$. The proof of this hardness result is given in Section 7. Therefore our aim is to design efficient approximation algorithms. We shall first consider the two-dimensional version of the problem and give a $\frac{3}{2}$-approximation algorithm for the minimum size globally rigid spanning subgraph problem as well as a 2-approximation algorithm for the metric version. We also show how the latter factor can be improved to 1.61 when the costs are defined by Euclidean distances in the plane.

In the second part of the paper we design constant factor approximation algorithms for the d-dimensional problem in the metric case, for all $d \geq 3$.

We can define – and we shall also consider – similar optimization problems by replacing global rigidity with redundant rigidity or rigidity (defined in Section 2 below) in the definition of MCGRSS. These problems are denoted by MCRRSS and MCRSS, respectively. It turns out that MCRRSS in $\mathbb{R}^d$ is also NP-hard for all $d \geq 1$. On the other hand, MCRSS is solvable in polynomial time in $\mathbb{R}^1$ and $\mathbb{R}^2$. The complexity status of MCRSS is open in $\mathbb{R}^d$ for $d \geq 3$.

1.2 Previous work

It is a well-known folklore result in rigidity theory that a graph G is redundantly rigid (resp. globally rigid) in $\mathbb{R}^1$ if and only if it is 2-edge-connected (resp. 2-connected). Thus in the one-dimensional case of MCRRSS (resp. MCGRSS) we search for a minimum cost 2-edge-connected (resp. 2-connected) spanning subgraph. These problems, even in the uniform or metric version, are NP-hard, as they contain the Hamilton cycle problem as a special case. There are several constant factor approximation algorithms in the literature that deal with these problems, see e.g. [23]. In light of this connection the MCGRSS problem is a natural extension of these core problems from graph connectivity.

The only higher dimensional result we are aware of is due to García and Tejel [11]. They consider the minimum size redundantly rigid augmentation problem in the plane, which corresponds to MCRRSS in the special case when G is complete and $c(e) \in \{0, 1\}$ for all $e \in E(G)$. They show that this problem is NP-hard in general but can be solved in polynomial time if the graph to be augmented - that is, the graph of the edges of cost zero - is minimally rigid in $\mathbb{R}^2$. The minimum size globally rigid augmentation problem is briefly mentioned in [9, 18], along with some related results.

1.3 Motivation and applications

One of the applications that inspired our research is the localization problem of two- and three-dimensional wireless sensor networks. In this problem the goal is to compute the locations of all sensors, when only a subset of the pairwise distances and locations is available. The network is localizable (that is, the localization problem has a unique solution) if and only if the corresponding framework is globally rigid [3]. In this framework the vertices correspond to the sensors and two vertices are adjacent if and only if the distance between them is known. Methods and results from rigidity theory have been used to solve a number of related problems. In particular, the characterization of localizability (assuming generic locations in the plane) and inductive constructions of localizable networks have been identified, see e.g. [1, 16]. Similar questions (concerning global rigidity or redundant rigidity) arise in molecular conformation, where the shape of a molecule is to be determined based on a subset of inter-atomic distances [30], in formation control [31], and elsewhere.

The minimum cost globally (or redundantly) rigid spanning subgraph problem may emerge in these applications when one wants to achieve, say, global rigidity by measuring (or recomputing, fixing, etc.) some pairwise distances in an optimal way. For example, it may happen that (i) certain distances are not computable, or more generally, the cost or time of computing pairwise distances may be different for different pairs, or preferences may be given to some pairs, or (ii) the level of noise in the distance data may be different, or (iii) the total length of the edges is a relevant factor, etc. These properties and parameters may be encodable in the cost and objective functions and then, assuming that the costs are uniform or metric, a near optimal solution can be obtained by using the approximation algorithms designed in this paper.

2 Rigid and globally rigid graphs

In this section we collect the basic definitions and results from rigidity theory that we shall use. The framework (G, p) is *rigid* in $\mathbb{R}^d$ if there exists an $\epsilon > 0$ such that, if (G, q) is equivalent to (G, p) and $||p(v) - q(v)|| < \epsilon$ for all $v \in V$, then (G, q) is congruent to (G, p). It is known that, informally speaking, this is equivalent to saying that every continuous motion of the vertices of the framework in $\mathbb{R}^d$ which preserves all edge-lengths takes the framework to a congruent realization of G. It is clear that global rigidity implies rigidity.

As for global rigidity, the rigidity of frameworks in $\mathbb{R}^d$ is a generic property for all $d \geq 1$ [2]. We say that a graph G is *rigid* in $\mathbb{R}^d$ if every (or equivalently, if some) generic realization of G in $\mathbb{R}^d$ is rigid. See Figure 1 for examples. A rigid graph $G = (V, E)$ in $\mathbb{R}^d$ is called *minimally rigid* if $G - e$ is not rigid for all $e \in E$.

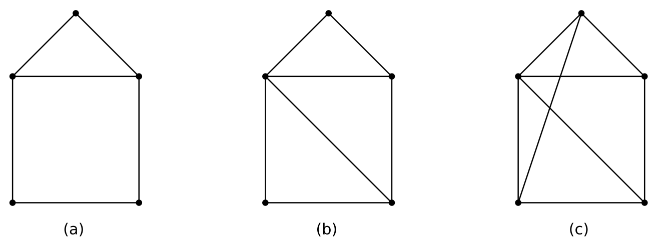

Figure 1: Graphs which are (a) not rigid, (b) rigid but not globally rigid, (c) globally rigid in the plane.

It is known that the edge sets of the minimally rigid graphs on vertex set V correspond to the bases of the so-called d-dimensional *rigidity matroid*, defined on the edge set of a complete graph on V. Hence they have the same number of edges: for example, a minimally rigid graph in $\mathbb{R}^2$ on vertex set V has $2|V| - 3$ edges. The problem of finding a polynomially verifiable characterization for graphs which are rigid in $\mathbb{R}^d$ has been solved for $d = 1, 2$, but is a major open problem for $d \geq 3$. We refer the reader to [19, 20] for more details on rigid and globally rigid frameworks and graphs.

In the plane we have the following key result. Let $G = (V, E)$ be a graph. For a subset $X \subseteq V$ we use $i(X)$ to denote the number of edges induced by X. We say that G is *sparse* if

$$i(X) \leq 2|X| - 3 \text{ for all } X \subseteq V \text{ with } |X| \geq 2. \tag{1}$$

The operation 0-*extension* adds a new vertex v to G and two new edges vx, vy for two distinct vertices x and y of G. The *1-extension* operation on edge uw and vertex z with $z \notin \{u, w\}$ adds a new vertex v, deletes uw, and adds three new edges vu, vw, vz. See Figure 2.

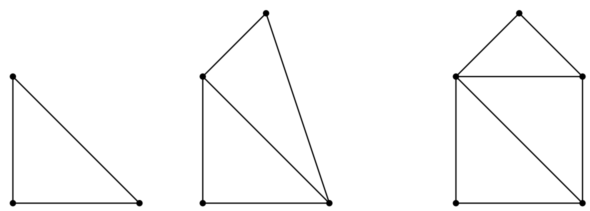

Figure 2: The graphs obtained from K_3 (left) by a 0-extension operation (middle) followed by a 1-extension operation (right).

The characterization of (minimally) rigid graphs is due to Laman.

Theorem 2.1. *[24] Let $G = (V, E)$ be a graph with $|E| = 2|V| - 3$. Then the following are equivalent:*
(i) G is minimally rigid in $\mathbb{R}^2$,
(ii) G is sparse,
(iii) G can be obtained from K_2 by a sequence of 0-extensions and 1-extensions.

As far as global rigidity is concerned, Hendrickson found the following necessary conditions for global rigidity in $\mathbb{R}^d$. We call a graph G *redundantly rigid* in $\mathbb{R}^d$ if $G - e$ is rigid in $\mathbb{R}^d$ for all $e \in E(G)$. A graph G is said to be k*-connected* if $G - X$ is connected for all $X \subset V(G)$ with $|X| \leq k - 1$.

Theorem 2.2. *[13] Let G be a globally rigid graph in $\mathbb{R}^d$ on at least $d + 2$ vertices. Then G is*
(i) $(d + 1)$-connected, and
(ii) redundantly rigid in $\mathbb{R}^d$.

These conditions together are also sufficient in $\mathbb{R}^1$ and $\mathbb{R}^2$. The one-dimensional result is folklore, see [14] for a proof. In the plane we have the following characterization.

Theorem 2.3. *[15] Let $G = (V, E)$ be a graph on at least four vertices. Then the following are equivalent:*
(i) G is globally rigid in $\mathbb{R}^2$,
(ii) G is 3-connected and redundantly rigid in $\mathbb{R}^2$,
(iii) G can be obtained from K_4 by a sequence of 1-extensions and edge additions.

We shall also use the following result of Nash-Williams [26]. Note that the graphs in the next theorem may have multiple edges.

Theorem 2.4. *[26] Let $G = (V, E)$ be a graph and let k be a positive integer. Then the edge set of G can be partitioned into k forests if and only if $i(X) \leq k|X| - k$ holds for all non-empty vertex sets $X \subseteq V$.*

2.1 Algorithms

The structural results presented in this section give rise to efficient combinatorial algorithms for testing whether a given graph $G = (V, E)$ is rigid, redundantly rigid, or globally rigid in the plane. These algorithms use the fact that the edge sets of the sparse subgraphs of a graph form the independent sets of the 2-dimensional rigidity matroid and boil down to the existence of an efficient subroutine for checking whether a graph is sparse or not. The matroidal property makes it possible to find a minimum cost rigid spanning subgraph of a rigid graph in $\mathbb{R}^2$ with respect to an arbitrary cost function on the edge set, in polynomial time. Each of these basic problems can be solved in $O(|V|^3)$ time or faster, see e.g. [4] for more details.

3 Minimum size globally rigid spanning subgraphs

In this section we present two simple approximation algorithms for the minimum size globally (resp. redundantly) rigid spanning subgraph problems. We show that if we delete edges as long as possible, in a greedy fashion, maintaining the global (or redundant) rigidity of the graph, then we end up with a close-to-optimal solution.

A graph $G = (V, E)$ is called *minimally globally (resp. redundantly) rigid* in $\mathbb{R}^2$ if it is globally (resp. redundantly) rigid in $\mathbb{R}^2$ but $G - e$ is not globally (resp. redundantly) rigid in $\mathbb{R}^2$ for all $e \in E$.

Theorem 3.1. *Suppose that $G = (V, E)$ is minimally globally rigid in $\mathbb{R}^2$ with $|V| \geq 4$. Then $|E| \leq 3|V| - 6$.*

Proof. Consider a sequence of graphs $G_1, G_2, ..., G_t$ for which $G_1 = K_4$, $G_t = G$, and G_i is obtained from G_{i-1} by an edge addition or 1-extension for all $2 \leq i \leq t$. Such a sequence exists by Theorem 2.3. Since G is minimally globally rigid, every edge addition operation used in this sequence adds an edge which will be split into two edges later by a 1-extension operation. This leads to a pairing, that is, a bijection between the added edges and a subset of the 1-extension operations. Each pair increases the number of vertices by one and the number of edges by three. A 1-extension operation alone increases the number of vertices by one and the number of edges by two. Thus, since K_4 satisfies $|E(K_4)| = 3|V(K_4)| - 6$, and the total number of edges added by the operations is not more than three times the number of added vertices, $G_t = G$ satisfies $|E| \leq 3|V| - 6$, as required. □

A globally (or redundantly) rigid graph G in $\mathbb{R}^2$ on vertex set V has at least $2|V| - 2$ edges by Theorems 2.2 and 2.1. Since testing global rigidity can be done in polynomial time, Theorem 3.1 leads to an efficient constant factor approximation algorithm.

Theorem 3.2. *There is a polynomial time $\frac{3}{2}$-approximation algorithm for the minimum size globally rigid spanning subgraph problem in $\mathbb{R}^2$.*

A similar situation holds for redundant rigidity. Here we use the following result (whose proof is substantially more complicated than that of Theorem 3.1).

Theorem 3.3. *[19] Suppose that $G = (V, E)$ is minimally redundantly rigid in $\mathbb{R}^2$ with $|V| \geq 7$. Then $|E| \leq 3|V| - 9$.*

As a corollary, we obtain:

Theorem 3.4. *There is a polynomial time $\frac{3}{2}$-approximation algorithm for the minimum size redundantly rigid spanning subgraph problem in $\mathbb{R}^2$.*

4 Structural properties of minimally rigid graphs

In the next two sections we consider the metric versions of the two-dimensional MCRRSS and MCGRSS problems. Our algorithms will first identify a minimum cost (minimally) rigid spanning subgraph of the input graph and then extend it to a feasible solution by adding new edges. In order to keep the total cost of these added edges low we need structural results on the minimally rigid subgraphs of a minimally rigid graph. We shall rely on some results of García and Tejel from [11] and also prove a number of new properties. In what follows a minimally rigid graph in the plane will be called a *Laman graph*.

4.1 The extreme classes of a Laman graph

Let $G = (V, E)$ be a rigid graph. We say that an edge $e \in E$ is *redundant* in G if $G - e$ is rigid. Thus G is redundantly rigid if every edge of G is redundant. As we noted above, the Laman graphs on vertex set V are the bases of the two-dimensional rigidity matroid defined on the edge set of a complete graph on V. In particular, if G is Laman then $G + e$ has a unique (matroid) circuit, the *fundamental circuit of e* with respect to G. From this viewpoint the next lemma easily follows from some basic properties of matroids.

Lemma 4.1. *Let $G = (V, E)$ be a Laman graph and let $e = ij$ be an edge for some $i, j \in V$. Then*
(i) There is a unique fundamental circuit in $G+e$, denoted by $C(ij)$ or $C(e)$. This circuit contains e. $(V(C(e)), E(C(e)) - e)$ is a Laman subgraph of G, denoted by $L(ij) = (V(ij), E(ij))$ or simply $L(e)$.
(ii) For every edge $e' \in E(ij)$ the graph $(V, E + e - e')$ is a Laman graph, in which the fundamental circuit of e' is $C(ij)$. Moreover, if $e' \notin E(ij)$ then $(V, E + e - e')$ is not a Laman graph,
(iii) If G' is a Laman subgraph of G with $\{i, j\} \subseteq V(G')$ then $L(ij)$ is a subgraph of G'. Thus $L(ij)$ is equal to the intersection of all Laman subgraphs L_h of G with $\{i, j\} \subseteq V(L_h)$.

In other words $E(ij)$ is equal to the set of edges of G that become redundant in $G + e$. We may define $L(ij)$ even if $ij \in E(G)$. In this case $L(ij)$ is the single edge ij and $C(e)$ is a graph consisting of two parallel copies of ij.

For every $i, j \in V(G)$ we say that $L(ij)$ is a *generated* Laman subgraph of G whose *generator* is the edge ij. A Laman graph G is called *narrow* if $G = L(ij)$ for some i, j, that is, if it can be made redundantly rigid by adding one new edge. See Figure 3. Otherwise it is said to be *wide*. We note that the authors in [11] use generated and non-generated, respectively, instead of narrow and wide. We feel the new terminology makes the statements and proofs more transparent.

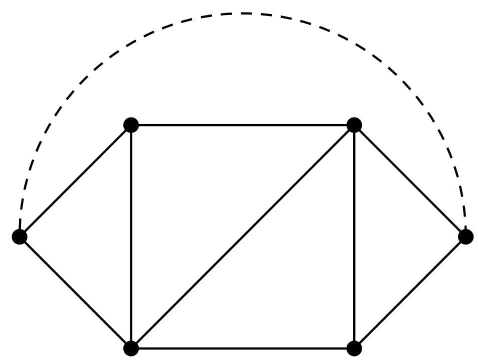

Figure 3: A narrow Laman graph (solid edges). Adding the dotted edge makes it redundantly rigid.

Given a Laman graph G and a set $e_1, e_2, ..., e_k$ of new edges, let $L(e_1, e_2, ..., e_k)$ be the subgraph of G consisting of those edges of G that are redundant in $G + \{e_1, e_2, ..., e_k\}$.

Lemma 4.2. *[11, Lemma 4] Let G be a Laman graph. Then* $L(e_1, e_2, ..., e_k) = L(e_1) \cup L(e_2) \cup ... \cup L(e_k)$.

Thus adding a set of new edges $e_1, e_2, ..., e_k$ to a Laman graph G yields a redundantly rigid graph if and only if the union of the fundamental circuits of the edges e_i, $1 \leq i \leq k$, contains every edge of G. In a smallest redundantly rigid augmentation of G we may assume that for every new edge e_i the fundamental circuit of e_i is a *maximal (with respect to inclusion) generated Laman subgraph*, or simply an *MGL*. A vertex i of G is said to be *extreme* if there is a vertex j for which $L(ij)$ is an MGL of G.

Let $G = (V, E)$ be a Laman graph and let X be the set of its extreme vertices. We say that $i, i' \in X$ are *equivalent* if there exists a vertex $j \in X$ for which $L(ij)$ is an MGL and $L(ij) = L(i'j)$. García and Tejel verified that this is an equivalence relation on X, assuming that G is wide [11, Lemma 8]. We call the equivalence classes of X defined by this relation the *extreme classes* of G. See Figure 4. The extreme vertices satisfy the following properties:

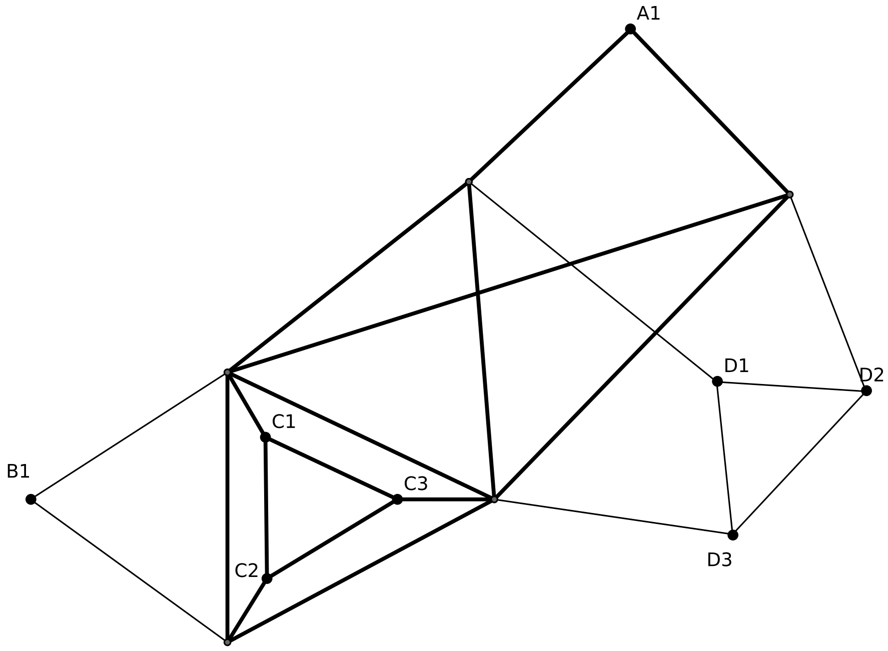

Figure 4: The extreme classes in a Laman graph. This graph has four extreme classes: $\{A_1\}$, $\{B_1\}$, $\{C_1, C_2, C_3\}$, and $\{D_1, D_2.D_3\}$. The edges of the MGL subgraph generated by an edge connecting A_1 to some C_i $(1 \leq i \leq 3)$ are thick.

Lemma 4.3. *[11, Lemma 9] Let G be a wide Laman graph and let i_1, i_2 be extreme vertices of G. Then*
(i) if i_1 and i_2 are not equivalent then $L(i_1 i_2)$ is an MGL,
(ii) if i_1 and i_2 are equivalent then $L(i_1 i_2)$ is not an MGL,
(iii) if L' is MGL then L' contains extreme vertices from exactly two extreme classes of G. If i_1, i_2 are vertices from these two classes then $L' = L(i_1 i_2)$.

The next result gives rise to an edge set whose addition makes every edge redundant.

Lemma 4.4. *[11, Lemma 10] Let G be a wide Laman graph. Suppose that G has h extreme classes with representative vertices $i_1, i_2, ..., i_h$. Then $G = \bigcup_{r=2}^{h} L(i_1 i_r)$.*

Thus G can be made redundantly rigid by adding $h - 1$ well chosen edges, based on the extreme classes. A more detailed analysis in [11] shows that in fact the optimum – the size of a smallest augmenting set – is equal to $\lceil \frac{h}{2} \rceil$, and that a set of representative vertices from the extreme classes as well as an optimal solution can be found in $O(n^2)$ time. We shall not use these facts concerning optimal augmentations but will rely on, and extend, some of the structural results on extreme classes from [11]. We shall use the following lemmas. The first one is well-known, see e.g. [15, Lemma 2.3].

Lemma 4.5. *Let $G = (V, E)$ be a Laman graph and let L_1, L_2 be Laman subgraphs of G with at least two vertices in common. Then their union as well as their intersection are also Laman subgraphs of G.*

Lemma 4.6. *[11, Lemma 5] Let G be a Laman graph on at least four vertices and let $L(ij)$ be an MGL subgraph of G. Then for every vertex $k \neq j$ the subgraphs $L(ij)$ and $L(jk)$ have at least one edge in common. In particular, $L(ij)$ contains all edges incident with i or j.*

A simple corollary is as follows.

Lemma 4.7. *[11] Let i, j, k be extreme vertices chosen from three different extreme classes. Then $L(ik) \subset L(ij) \cup L(jk)$.*

Proof. Since $L(ij)$ and $L(jk)$ are MGL subgraphs of G, Lemma 4.6 implies that every edge incident with j belongs to both. Thus they have at least two vertices in common, which gives, by Lemma 4.5, that $L(ij) \cup L(jk)$ is a Laman subgraph of G. As it contains i and j, we must have $L(ik) \subset L(ij) \cup L(jk)$ by Lemma 4.1(iii). □

4.2 Extreme classes and separating pairs

Since the globally rigid graphs in the plane are 3-connected, a new set of edges whose addition to a Laman graph makes it globally rigid must eliminate all separating pairs. In order to handle this condition we next prove new structural results on the relation between extreme classes and separating pairs. We start with two preliminary lemmas about wide Laman graphs.

Lemma 4.8. *Let $i_1, i_2, .., i_q$ be extreme vertices chosen from q different extreme classes. Then $L(i_1 i_q) \subset L(i_1 i_2) \cup L(i_2 i_3) \cdots \cup L(i_{q-1} i_q)$.*

Proof. We apply induction on q. For $q = 3$ the lemma follows from Lemma 4.7. Now suppose that $q \geq 4$ and the lemma holds up to $q - 1$. Then $L(i_1 i_q) \subset L(i_1 i_{q-1}) \cup L(i_{q-1} i_q) \subset L(i_1 i_2) \cup \cdots \cup L(i_{q-2} i_{q-1}) \cup L(i_{q-1} i_q)$. □

Lemma 4.9. *Let T be a set of extreme vertices of G that contains exactly one vertex from each extreme class and let F be a set of edges for which (T, F) is connected. Then $G + F$ is redundantly rigid.*

Proof. It follows from Lemma 4.4 that there exists an edge set J for which every edge of J is induced by T and $G + J$ is redundantly rigid. By Lemma 4.8 and the connectivity of (T, F) it follows that $G + F$ is redundantly rigid. □

Let $G = (V, E)$ be a 2-connected graph. We say that a pair $\{u, v\} \subset V$ is a *separating pair* in G if $G - \{u, v\}$ is disconnected. If X is the vertex set of a connected component of $G - \{u, v\}$, for some separating pair $\{u, v\}$, then X is called a *fragment*. For a vertex set $Z \subseteq V$ a vertex $w \in V - Z$ is called a *neighbour* of Z if there is an edge from w to some vertex of Z. The set of neighbours of Z is denoted by $N(Z)$. Thus $N(X)$ forms a separating pair for every fragment X. A minimal fragment of G (with respect to inclusion) is an *end*. A separating pair $\{u_1, v_1\}$ *crosses* another separating pair $\{u_2, v_2\}$ if u_1 and v_1 belong to different components of $G - \{u_2, v_2\}$. It is not hard to see that if $\{u_1, v_1\}$ crosses $\{u_2, v_2\}$ then $\{u_2, v_2\}$ crosses $\{u_1, v_1\}$. Hence these pairs are said to be *crossing* separating pairs. The next lemma is easy to verify.

Lemma 4.10. *Let G be 2-connected and suppose that v is a vertex of some end B of G. If $\{u, v\}$ is a separating pair for some vertex u then $N(B)$ and $\{u, v\}$ are crossing separating pairs.*

Lemma 4.11. *[15, Lemmas 2.6(a), 3.5(b)] Let G be a rigid graph on at least three vertices. Then*
(i) G is 2-connected, and
(ii) there are no crossing separating pairs in G.

Given two disjoint vertex sets $X, Y \subseteq V$ in a graph, the number of edges from X to Y is denoted by $d(X, Y)$.

Lemma 4.12. *Let $G = (V, E)$ be a Laman graph and let $X, Y \subset V$ with $|X \cap Y| = \{u, v\}$ and $d(X - Y, Y - X) = 0$. Then*
(i) if $uv \in E$ and $V = X \cup Y$ then $G[X]$ and $G[Y]$ are both Laman,
(ii) if $uv \notin E$ then at most one of $G[X]$ and $G[Y]$ is Laman. Furthermore, if $V = X \cup Y$ then exactly one of $G[X]$ and $G[Y]$ is Laman.

Proof. First suppose that $uv \in E$ and $V = X \cup Y$. Then we have $2|V| - 3 = i(X) + i(Y) - 1 \leq 2|X| - 3 + 2|Y| - 3 - 1 = 2|V| - 3$. This implies (i). Next suppose that $uv \notin E$. If $G[X]$ and $G[Y]$ are both Laman then we have $2|X \cup Y| - 3 \geq i(X \cup Y) = i(X) + i(Y) = 2|X| - 3 + 2|Y| - 3 = 2|X \cup Y| - 2$, a contradiction. This proves the first part of (ii). By assuming that $V = X \cup Y$ and that neither of $G[X]$ or $G[Y]$ is Laman we have $2|V| - 3 = i(X) + i(Y) \leq 2|X| - 4 + 2|Y| - 4 = 2|V| - 4$, a contradiction. This completes (ii). □

For a separating pair $\{u, v\}$ and a component C of $G - \{u, v\}$ let $\bar{C} = G[V(C) \cup \{u, v\}]$ be its *closure*.

Lemma 4.13. *Let G be a Laman graph and $\{u, v\}$ be a separating pair in G. Let the components of $G - \{u, v\}$ be denoted by $C_1, C_2, \ldots, C_t$. Then*
(i) if $uv \in E$ then $\bar{C}_i$ is Laman for all $1 \leq i \leq t$,
(ii) if $uv \notin E$ then there is a unique component, say C_1, for which $\bar{C}_1$ is Laman,
(iii) if $uv \notin E$ then $L(uv)$ intersects exactly one component of $G - \{u, v\}$.

Proof. First observe that (i) follows by applying Lemma 4.12(i) to the sets $X = V(\bar{C}_i)$ and $Y = V - V(C_i)$. Next we assume $uv \notin E$. Then Lemma 4.12(ii) gives that at most one $\bar{C}_i$ is Laman. For a contradiction suppose that no $\bar{C}_i$ is Laman. Then $2|V| - 3 = |E| = \sum_1^t i(\bar{C}_j) = \sum_1^t (2|V(\bar{C}_j)| - 4) = 2|V| + 4(t-1) - 4t = 2|V| - 4$, a contradiction. Finally, (iii) follows from (ii), since if $\bar{C}_1$ is Laman then it must contain the unique smallest Laman subgraph $L(uv)$ containing u, v. □

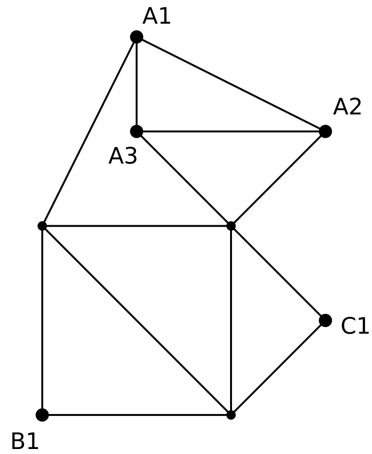

Figure 5: In this graph A_1 is an extreme vertex that belongs to a separating pair.

In a Laman graph an extreme vertex may belong to some separating pair, see Figure 5. The next lemmas will show that it cannot happen to all vertices of an extreme class.

Lemma 4.14. *Let G be a Laman graph and let $\{u, v\}$ be a separating pair in G. Consider a pair x, y of vertices with $x \in A, y \in B$, where A, B are distinct connected components of $G - \{u, v\}$. Then $L(uv) \subseteq L(xy)$ and $L(ux) \subseteq L(xy)$.*

Proof. By Lemma 4.1(iii), $L(xy)$ is a (smallest) Laman subgraph that contains x and y. Since $xy \notin E$, $L(xy)$ is 2-connected by Lemma 4.11(i). Thus we have $\{u, v\} \subseteq L(xy)$. Hence $L(uv) \subseteq L(xy)$. A similar argument gives $L(ux) \subseteq L(xy)$. □

Lemma 4.15. *Suppose that G is a wide Laman graph. Then every extreme class of G contains at least one vertex which is not part of any separating pair in G.*

Proof. Consider an extreme class P of G and fix an extreme vertex $u \in P$. Suppose that $\{u, v\}$ is a separating pair for some $v \in V$. Since u is extreme, there exists an MGL $L(uj)$ for some extreme vertex j. Fix two components A, B of $G - \{u, v\}$ and a pair of vertices $x \in A, y \in B$.

We claim that $j \neq v$. To see this first note, that if $uv \in E$ then $L(uv)$ is not an MGL, and hence $j \neq v$ follows. Next suppose that $uv \notin E$. Then we have $L(uv) \subseteq L(xy)$ by Lemma 4.14. Furthermore, the inclusion must be proper by Lemma 4.13(iii). This shows that that $L(uv)$ is not an MGL. Hence $j \neq v$ and the claim follows.

By symmetry we may assume that $j \notin B$. Then it follows from Lemma 4.14 that for every vertex $y \in B$ we have $L(uj) = L(yj)$ and hence y is also in P. By taking y to be a vertex of some end within B the lemma follows from Lemmas 4.10 and 4.11(ii). □

A similar result holds for narrow Laman graphs.

Lemma 4.16. *Let $G = (V, E)$ be a narrow Laman graph. Then there is a pair $u, v \in V$ which is disjoint from all separating pairs and for which $G + uv$ is redundantly rigid.*

Proof. Since G is narrow, there is a pair $u_1, v_1 \in V$ for which $G + u_1v_1$ is redundantly rigid. Suppose that $\{u_1, w\}$ is a separating pair for some vertex $w \in V$. By Lemma 4.13(iii) we must have $w \neq v_1$. Let A, B be two connected components of $G - \{u_1, w\}$ with $v_1 \in A$ and consider a vertex $u \in B$. By Lemma 4.14 we have $L(u_1v_1) \subseteq L(uv_1)$. Since $L(u_1v_1) = G$, it follows that $G + uv_1$ is redundantly rigid. $L(xv_1) = G$. By choosing u to be a vertex of some end within B, we may assume that u is disjoint from all separating pairs. Now applying a similar argument to the pair $\{u, v_1\}$ we obtain that there is a pair $\{u, v\}$ which is disjoint from all separating pairs and for which $L(uv) = G$. This completes the proof. □

5 Minimum cost globally rigid spanning subgraphs

In this section we consider the metric MCRRSS and MCGRSS problems in the plane. To illustrate the main ideas, we start with the minimum cost redundantly rigid subgraph problem, for which we have a simpler approximation algorithm. Recall that the input of both problems is a complete graph $K = (V, E(K))$ on at least four vertices and a metric cost function $c : E(K) \to \mathbb{R}_+$.

Algorithm MinCostRedRig2

(i) Compute a minimum cost spanning Laman subgraph $G = (V, E)$ of K.

(ii) If G is a wide Laman graph then find a set S of extreme vertices of G that contains exactly one vertex from each extreme class and compute a minimum cost spanning tree (S, F) of $K[S]$, where $K[S]$ is the subgraph of K induced by S. Output $(V, E + F)$.

(iii) If G is a narrow Laman graph then find a new edge e for which $G + e$ is redundantly rigid. Output $(V, E + e)$.

Theorem 5.1. *Algorithm MinCostRedRig2 is a polynomial time 2-approximation algorithm for the metric MCRRSS in $\mathbb{R}^2$.*

Proof. Consider an instance of MCRRSS. If G is wide, the output is a feasible solution by Lemma 4.9. If G is narrow, the output is feasible by construction. To verify the approximation ratio consider an optimal solution G^*. Let OPT denote the total cost of the edges of G^*. Since G^* is rigid, we have $c(E) \leq OPT$. We claim that G^* contains two edge-disjoint spanning trees. Indeed, since G^* is redundantly rigid, there exists a minimally rigid spanning subgraph H of $G^* - e$, for any fixed edge e of G^*: now Theorems 2.1 and 2.4 imply that $H + e$ is the union of two edge-disjoint spanning trees.

Suppose that G is wide and output is obtained in step (ii). Since G^* contains two edge-disjoint spanning trees, a minimum cost spanning tree F^* of K satisfies $c(F^*) \leq \frac{OPT}{2}$. Furthermore, it is well-known that if c is metric and $S \subseteq V(G)$ then the cost of a minimum cost spanning tree in $K[S]$ has cost at most $2c(F^*)$. This follows by doubling the edges of F^* to obtain an Eulerian graph J and then shortcutting an Eulerian walk of J to obtain a spanning cycle C on S. Since c is metric and C contains a spanning tree of $K[S]$, it follows that the minimum cost spanning tree on S has cost at most $2c(F^*)$. Hence if G is wide then we have $c(E + F) \leq 2OPT$, as required.

Next suppose that G is narrow and the output is obtained in step (iii). Let $e = uv$ be the edge found for which $G + e$ is redundantly rigid. Since G^* contains two edge-disjoint uv-paths, there is a uv-path P with $c(P) \leq \frac{OPT}{2}$. By using that c is metric, we obtain $c(e) \leq c(P) \leq \frac{OPT}{2}$ and hence $c(E + e) \leq \frac{3}{2}OPT \leq 2OPT$, as claimed.

The polynomial running time of the algorithm follows by noting that a minimum cost spanning tree or a minimum cost spanning Laman subgraph can be found efficiently by a greedy algorithm. Moreover, as we remarked earlier, the extreme classes of G can also be found in polynomial time. □

Next we consider the metric MCGRSS in $\mathbb{R}^2$. The following algorithm is a refined version of Algorithm MinCostRedRig2.

Algorithm MinCostGlobRig2

(i) Compute a minimum cost spanning Laman subgraph $G = (V, E)$ of K.

(ii) If G is a wide Laman graph then find a set S of extreme vertices of G that contains exactly one vertex from each extreme class, so that the vertex belongs to no separating pair of G.

(iii) If G is a narrow Laman graph then find a pair $S = \{i, j\}$ of vertices, for which $G + ij$ is redundantly rigid and i, j belong to no separating pair of G.

(iv) Find a set T of vertices of G that contains exactly one vertex from each end W of G which is disjoint from S,

(v) Compute a minimum cost spanning tree (R, F) of $K[R]$, where $R = S \cup T$. Output $(V, E + F)$.

The steps of the algorithm are well-defined by Lemmas 4.15 and 4.16. We next show that the output is a feasible solution.

Lemma 5.2. *The output of Algorithm MinCostGlobRig2 is*
(i) 3*-connected, and*
(ii) redundantly rigid.

Proof. First we prove (i). By the choice of the vertices in S (c. f. Lemmas 4.15, 4.16) and the vertices in T added from the ends (c.f. Lemmas 4.10, 4.11) no vertex in R belongs to a separating pair of G. Furthermore, for every end (and hence for every fragment) X we must have $X \cap R \neq \emptyset$. This implies that adding a tree on R eliminates every separating pair of G and hence makes it 3-connected.

Next we prove (ii) simultaneously for the two cases, that depend on whether G is wide or narrow. Let us fix two vertices $i, j \in S$ for which every internal vertex of the path P from i to j in (R, F) is a vertex in T. Let $P = i, t_1, t_2, \ldots, t_r, j$ and $L_i = L(t_i t_{i+1})$ for $1 \leq i \leq r - 1$. The key observation, which follows from Lemma 4.14, is that in the sequence $L(it_1)$, L_1, L_2, $\ldots L_{r-1}$, $L(t_r j)$ each pair of consecutive Laman subgraphs have at least two vertices in common. By Lemma 4.5 this implies that their union is Laman. Hence $L(ij) \subseteq L(it_1) \cup L(t_1t_2) \cup \ldots \cup L(t_{r-1}t_r) \cup L(t_rj)$. Then it follows that by adding the edges of P we make every edge of $L(ij)$ redundant. Therefore, by Lemma 4.9, adding F makes every edge of G redundant. $\square$

An analysis similar to that of MinCostRedRig2, together with Lemma 5.2 above, gives our main result in $\mathbb{R}^2$.

Theorem 5.3. *Algorithm MinCostGlobRig2 is a polynomial time* 2*-approximation algorithm for the metric MCGRSS in* $\mathbb{R}^2$.

We have a family of instances showing that the approximation ratio of Algorithm MinCostRedRig2 (and of MinCostGlobRig2) is not better than $\frac{3}{2}$. Consider a complete graph K on $2s+1$ vertices, for some integer $s \geq 2$. Fix a subset E of vertices of size s and define the costs of the edges of K so that the edges between vertices in E are of cost 2 while the cost of every other edge is equal to 1. The algorithm may find, as the minimum cost spanning Laman subgraph, a graph in which each vertex in E is an extreme vertex of degree two. See Figure 6 for the case $s = 5$. The minimum cost tree on these vertices has total cost $2s-2$. Thus the output has cost $4s - 1 + 2s - 2 = 6s - 3$. On the other hand it is not hard to see that a feasible solution of cost $4s$ exists.

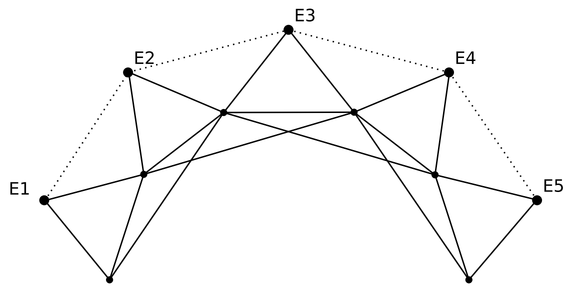

Figure 6: The solid edges correspond to the spanning Laman subgraph. The dotted edges form a tree on its extreme vertices.

5.1 The Euclidean case

In the Euclidean version of our problems the vertices correspond to points in $\mathbb{R}^2$ and the cost of an edge is the Euclidean distance of its endpoints. In this version, which may occur for example in the network localization problem, our algorithm has a better approximation ratio.

In order to show this, recall that in the Euclidean Steiner Tree Problem we are given a set S of points in the plane and the goal is to find a tree of minimum total length, which contains S. The tree may use points not in S. The ratio of the total length of a shortest spanning tree on S and the total length of a shortest Steiner tree with respect to S is the so called *Steiner ratio*. It was proved in [5] that the Steiner ratio is at most 1.22.

We can use this fact in the analysis of our algorithm and deduce that $c(F) \leq 1.22c(F^*) \leq 0.61OPT$, following the notation of Theorem 5.1. Thus the approximation ratio of the Euclidean version of MinCostRedRig2 (and MinCostGlobRig2) is 1.61.

6 Higher dimensions

In this section we design an approximation algorithm for the d-dimensional metric MCGRSS problem, which works for every $d \geq 2$, with an approximation ratio that depends only on d.

The algorithm is rather simple and is based on the idea of graph powers. The k^{th} *power* of graph G, denoted by G^k, is the graph on the same vertex set, in which two vertices are adjacent if and only if their distance in G is at most k. The input of the algorithm is an integer $d \geq 2$, a complete graph $K = (V, E(K))$ on at least $d + 2$ vertices and a metric cost function $c : E(K) \to \mathbb{R}_+$.

Algorithm MinCostGlobRigGen

(i) Compute a minimum cost spanning tree T of K.
(ii) By shortcutting $2T$ create a Hamilton cycle C on vertex set V.
(iii) Output C^d.

In step (ii) the graph $2T$ is obtained from T by replacing every edge of T by two parallel edges. The shortcutting operation is standard and we already used in the analysis of Theorem 5.1: we find an Eulerian walk of $2T$ and by shortcutting repeated vertices we turn it into a Hamilton cycle.

The fact that the output is a feasible solution follows from the next lemma. Let C_n denote a cycle on n vertices.

Lemma 6.1. *C_n^d is globally rigid in $\mathbb{R}^d$.*

Proof. If $n \leq 2d + 1$ then C_n^d is complete, and hence globally rigid in $\mathbb{R}^d$. So we may assume that $n \geq 2d + 2 \geq d + 2$. We shall prove that a spanning subgraph of C_n^d can be obtained from K_{d+2}, which is globally rigid, by a sequence of (d-dimensional) 1-extensions. This operation adds a new vertex v to the graph, deletes an edge uw, and adds $d + 1$ new edges incident with v, so that the set of new edges includes vu and vw. It is known that this operation preserves global rigidity in $\mathbb{R}^d$, see [6].

Label the vertices of C_n^d by $v_1, ..., v_n$ and start with a K_{d+2} on vertex set $v_1, .., v_{d+2}$. In the first iteration perform a 1-extension which adds vertex v_{d+3}, deletes the edge $v_1 v_{d+2}$, and connects v_{d+3} to $v_{d+2}, v_{d+1}, .., v_3$ and v_1. In the next iteration add v_{d+4} by a 1-extension on edge $v_1 v_{d+3}$ so that the new vertex is connected to the preceding d vertices and to v_1, and so on. After $n - d - 2$ iterations all vertices of C_n^d are included and the graph constructed is a globally rigid spanning subgraph of C_n^d. See Figure 7. This completes the proof. □

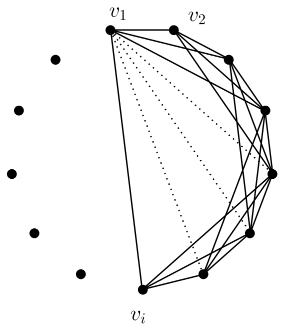

Figure 7: A graph obtained in the process of constructing a globally rigid spanning subgraph of C_{13}^3 by 1-extensions. The last vertex added up to this point is v_i. The dotted edges have been deleted by the previous 1-extensions.

The analysis of the algorithm will also use the following claim.

Lemma 6.2. *Suppose that $G = (V, E)$ is rigid in $\mathbb{R}^d$ with $|E| = d|V| - d$ and $|V| \geq d + 1$, for some $d \geq 1$. Then the edge set of G can be decomposed into d spanning trees.*

Proof. Since $|E| = d(|V| - 1)$, it suffices to show that the edge set of G can be partitioned into d forests. We shall verify that G satisfies the condition in Theorem 2.4 for $k = d$. Before counting edges let us fix a minimally rigid spanning subgraph H of G. It is known that H has $d|V| - \binom{d+1}{2}$ edges and for every vertex set $X \subseteq V$ with $|X| \geq d + 1$ we have $i_H(X) \leq d|X| - \binom{d+1}{2}$. The number of edges of G which do not belong to H is equal to $\binom{d+1}{2} - d$.

Let $X \subseteq V$ be a non-empty vertex set. First suppose $|X| \geq d+1$. Then, by using the above bounds, we have $i_G(X) \leq i_H(X) + \binom{d+1}{2} - d \leq d|X| - d$, as required. Next suppose $|X| \leq d$. Then, since G has no parallel edges, we have $i_G(X) \leq \binom{|X|}{2} = \frac{|X|(|X|-1)}{2} \leq d(|X| - 1) = d|X| - d$. This completes the proof. □

We are ready to analyse the algorithm. For simplicity we shall assume that $|V| \geq \binom{d}{2}$. We remark that if the input graph is smaller, a similar analysis gives the upper bound $2d + 2$ for the approximation ratio. Moreover, in this case enumerating all feasible solutions would also be an option for d fixed.

Theorem 6.3. *Algorithm MinCostGlobRigGen is a polynomial time $(d+\frac{2d}{d-1})$-approximation algorithm for the metric MCGRSS problem in $\mathbb{R}^d$, assuming that the size of the input graph is at least $\binom{d}{2}$.*

Proof. The output is a feasible solution by Lemma 6.1. The polynomial running time is also clear. It remains to prove the approximation ratio.

Let $G^* = (V, E)$ be an optimal solution. Since it is globally rigid in $\mathbb{R}^d$, it is also redundantly rigid in $\mathbb{R}^d$ by Theorem 2.2. Thus we have $|E| \geq d|V| - \binom{d+1}{2} + 1$. Furthermore, $G^* = (V, E)$ is rigid in $\mathbb{R}^{d-1}$, too (an observation that follows easily from e.g. by the coning theorem of [7]). Also, since $|V| \geq \binom{d}{2}$, we have $|E| \geq d|V| - \binom{d+1}{2} + 1 = (d-1)|V| - (\binom{d}{2} +$

$d) + 1 + |V| \geq (d-1)|V| - (d-1)$. Thus we can apply Lemma 6.2 to G^* and deduce that it contains $d-1$ pairwise edge-disjoint spanning trees. Hence $c(F) \leq \frac{OPT}{d-1}$.

Therefore $c(C) \leq \frac{2OPT}{d-1}$. By using the metric property of c, the total cost of the edges of C^d that connect vertices which are of distance exactly k in C can be bounded by $kc(C)$. Thus

$$c(C^d) \leq (1+2+\cdots+d)c(C) = \frac{d(d+1)}{2}c(C)$$
$$\leq \frac{d(d+1)}{2}\frac{2}{d-1}OPT = \frac{d(d+1)}{d-1}OPT = \Big(d + \frac{2d}{d-1}\Big)OPT,$$

as claimed. □

Note that the approximation ratio of MinCostGlobRigGen for $d = 2$ (and for $d = 3$) is equal to 6, which is substantially worse than that of algorithm MinCostGlobRig2. In the next subsection we show how to improve on this ratio in the three-dimensional case by using a more sophisticated analysis.

6.1 Improving the ratio for $d = 3$

We start with a technical lemma.

Lemma 6.4. *Let $K = (V, E)$ be a complete graph and let $c : E \to \mathbb{R}_+$ be a metric cost function. Suppose that $G = (V, F)$ is a 3-connected spanning subgraph of K which contains no subgraph isomorphic to K_6. Then for every $p > 0$ there is an N_p such that if $|V| \geq N_p$ then there exists a pair $\{e, f\} \subset E - F$ of edges with $c(e) + c(f) \leq c(G)p$.*

Proof. First suppose that there is a vertex v with $d_G(v) \geq 3 + 4\lceil\frac{1}{p}\rceil$. Let $X = N_G(v)$. We claim that X induces at least $2\lceil\frac{1}{p}\rceil$ pairwise disjoint non-edges. Indeed, such a collection M of non-edges can be obtained in a greedy manner, using the fact that any subset of six vertices of X induces at least one non-edge. By using the metric property of c we can now deduce that

$$c(F) \geq \sum_{vu: u \in N_G(v)} c(vu) \geq \sum_{e \in M} c(e).$$

Thus the two edges e, f of M with the smallest cost satisfy $c(e) + c(f) \leq c(G)p$, as required.

Next suppose that $d_G(v) < 3 + 4\lceil\frac{1}{p}\rceil$ for all $v \in V$. Then, assuming $|V| > \sum_{i=0,\ldots,k}(3+4\lceil\frac{1}{p}\rceil)^i$ for some integer k, it follows that there exist two vertices $v_1, v_2 \in V$ for which the length of a shortest path from v_1 to v_2 in G is at least $k+1$.

Take three internally disjoint chordless paths from v_1 to v_2. Let P be one of them with minimum total cost. Then we have $c(P) \leq \frac{1}{3}c(G)$. Furthermore, by taking a path of non-edges connecting every second vertex along P and using the fact that c is metric we obtain a set N of at least $\lfloor\frac{k+1}{2}\rfloor$ non-edges with $c(N) \leq c(P) \leq \frac{1}{3}c(G)$. Thus the two edges e, f of N with the smallest cost satisfy $c(e) + c(f) \leq \frac{1}{3}\frac{2}{\lfloor\frac{k+1}{2}\rfloor}c(G) \leq \frac{1}{3k}c(G)$. Hence by choosing $k \geq \frac{1}{3p}$ and $k \geq 4$, we have $c(e) + c(f) \leq c(G)p$, as required. □

The next lemma leads to an improved bound by choosing arbitrary small $p < \frac{3}{2}$.

Lemma 6.5. *Let $K = (V, E)$ be a complete graph and let $c : E \to \mathbb{R}_+$ be a metric cost function. Suppose that $G = (V, F)$ is a globally rigid subgraph of K in $\mathbb{R}^3$ and let $H = (V, T)$ be a minimum cost spanning tree. Then for every $p > 0$ there is an N_p such that if $|V| \geq N_p$ then $c(T) \leq \frac{1}{3}(1 + p)c(G)$.*

Proof. We shall use that, since G is globally rigid in $\mathbb{R}^3$, G is 3-connected and has a spanning proper subgraph G' with $3|V| - 6$ edges satisfying $i_{G'}(X) \leq 3|X| - 6$ for all $X \subseteq V$ with $|X| \geq 3$. Let's fix p.

First suppose $|F| \geq 3|V| - 3$. Then we can add three edges from F to G' and obtain a subgraph of G which contains three edge-disjoint spanning trees (by Nash-Williams' theorem). Hence $c(T) \leq \frac{2}{3}c(G)$.

Next suppose $3|V| - 5 \leq |F| \leq 3|V| - 4$. Then the sparsity property of G' implies that G contains no subgraph isomorphic to K_6. Now we may apply Lemma 6.4 to G and deduce that there is an N_p such that if $|V| \geq N_p$ then there exists a pair $\{e, f\} \subset E - F$ of edges with $c(e) + c(f) \leq c(G)p$. A similar argument gives that $G + e + f$ contains three edge-disjoint spanning trees, and hence $c(T) \leq \frac{1}{3}(1 + p)c(G)$. □

We can now deduce an upper bound on the approximation ratio of MinCostGlobRigGen for $d = 3$, at least for large enough graphs, which can be arbitrarily close to 4.

Theorem 6.6. *Let $K = (V, E)$ be a complete graph and let $c : E \to \mathbb{R}_+$ be a metric cost function. For every $p > 0$ there is an N_p such that if $|V| \geq N_p$ then the approximation ratio of MinCostGlobRigGen for $d = 3$ is at most $2\frac{1}{3}(1 + p)6 = 4(1 + p)$.*

7 Concluding remarks

In this paper we introduced the Minimum cost globally rigid spanning subgraph problem in $\mathbb{R}^d$ and gave polynomial time approximation algorithms for the metric version. It remains an open problem to find similar results for general cost functions.

For Euclidean costs we obtained a somewhat better approximation ratio. It might be possible to find a polynomial time approximation scheme, like in the case of the k-connected spanning subgraph problem, see e.g. [8].

Finally we remark that a long list of similar problems can be obtained by replacing global rigidity in $\mathbb{R}^2$ (or equivalently, 3-connectivity and redundant rigidity) by other types of connectivity and sparsity requirements. The matroid on the edge set of a graph defined by the sparsity count of 1 happens to be a specific example of the so-called *count matroids*. These matroids can defined in a similar way by replacing $i(X) \leq 2|X| - 3$ by $i(X) \leq k|X| - l$ for some integers k, l with $l \leq 2k$, see [10, 29]. One can also define "redundant rigidity" with respect to these more general counts in a natural way. Partial results, extending the work in [11], have already been obtained by C. Király [22].

7.1 Hardness results

For completeness we show that the problems considered in this paper are NP-hard. Since global rigidity is equivalent to 2-connectivity in $\mathbb{R}^1$, finding a smallest globally rigid spanning subgraph of a graph G on the line is more general than the Hamilton cycle problem. Hence MCGRSS is NP-hard in $\mathbb{R}^1$. By applying a sequence of $d-1$ coning operations[1] to G, and assigning cost zero to each of the new edges, we can reduce the problem to the d-dimensional MCGRSS problem, for any given d, showing that MCGRSS is NP-hard in $\mathbb{R}^d$. A similar argument shows that MCRRSS is also NP-hard in $\mathbb{R}^d$ for all d.

A slightly more involved argument shows that these problems remain NP-hard in the metric case. Here we give the proof for MCGRSS in $\mathbb{R}^2$. Similar arguments can be used to extend the result to higher dimensions and to prove the hardness of metric MCRRSS in $\mathbb{R}^d$.

Theorem 7.1. *It is NP-hard to find a minimum cost globally rigid spanning subgraph in* $\mathbb{R}^2$ *of a given complete graph* $G = (V, E)$ *with respect to a metric cost function* $c : E \to \mathbb{R}$.

Proof. We shall reduce the Hamilton cycle problem to our problem. Consider an instance $H = (V, E)$ of the Hamilton cycle problem. Let G be the cone of H, where the new vertex is denoted by v, and let K be the complete graph on vertex set $V \cup \{v\}$. We assign costs to the edges of K as follows.

For every edge $e = uv$ with $u, v \in V$ we let $c(e) = 1.1$ (resp. $c(e) = 1.9$) if $uv \in E$ (resp. if $uv \notin E$). For the remaining edges e of K, which are incident with v, we define $c(e) = 1$. We claim that H has a Hamilton cycle if and only if the minimum cost globally rigid spanning subgraph of K, with respect to c, has total cost $2.1|V|$.

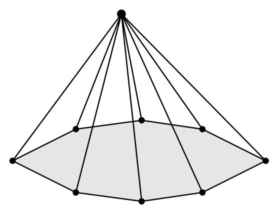

Figure 8: The cone graph of a graph.

To see this first suppose that H has a Hamilton cycle C. It is easy to see that the cone graph of C is globally rigid in $\mathbb{R}^2$. The total cost of the cone is $1.1|V| + |V| = 2.1|V|$. Next suppose that there is a globally rigid spanning subgraph F of K with cost at most $2.1|V|$. Since every globally rigid subgraph of K has at least $2|V|$ edges, the definition of c implies that F has exactly $2|V|$ edges and that it contains every edge incident with v.

Thus F is the cone graph of a 2-connected spanning subgraph C of H with $|V|$ edges. This shows that C is a Hamilton cycle in H. This completes the proof. □

Acknowledgement

This work was supported by the Hungarian Scientific Research Fund grant no. K109240 and K115483, and the ÚNKP-18-3 New National Excellence Program of the Ministry of

[1]The *cone* of graph G is obtained from G by adding a new vertex v and new edges from v to all vertices of G. See Figure 8. Connelly and Whiteley [7] proved that a graph G in globally rigid in $\mathbb{R}^d$ if and only if the cone of G is globally rigid in $\mathbb{R}^{d+1}$.

Human Capacities, Hungary.

References

[1] B. D. O. Anderson, P. N. Belhumeur, T. Eren, D. K. Goldenberg, A. S. Morse, W. Whiteley, and Y. R. Yang. Graphical properties of easily localizable sensor networks. *Wireless Networks* 15 (2): 177-191, 2009.

[2] L. Asimow and B. Roth. The rigidity of graphs, *Trans. Amer. Math. Soc.*, 245 (1978), pp. 279-289.

[3] J. Aspnes, T. Eren, D. K. Goldenberg, A. S. Morse, W. Whiteley, Y. R. Yang, B. D. O. Anderson, and P. N. Belhumeur. A theory of network localization. *IEEE Transactions on Mobile Computing* Vol. 5 (2006) 1663-1678.

[4] A.R. Berg and T. Jordán. Algorithms for graph rigidity and scene analysis. Proc. 11th Annual European Symposium on Algorithms (ESA) 2003, (G. Di Battista, U. Zwick, eds) Springer Lecture Notes in Computer Science 2832, pp. 78-89, 2003.

[5] F.R.K. Chung and R.L. Graham. A new bound for Euclidean Steiner minimal trees, *Annals of the New York Academy of Sciences*, 440 (1985), pp. 328-346.

[6] R. Connelly. Generic global rigidity, *Discrete Comput. Geom.* 33 (2005), pp 549-563.

[7] R. Connelly and W. Whiteley. Global rigidity: the effect of coning, *Discrete Comput. Geom.* (2010) 43: 717–735.

[8] A. Czumaj and A. Lingas. Approximation schemes for minimum-cost k-connectivity problems in geometric graphs, in: Handbook of approximation algorithms and metaheuristics, T.F. Gonzalez (ed.), CRC, 2007.

[9] Z. Fekete, T. Jordán, Uniquely localizable networks with few anchors, Proc. Algosensors 2006, (S. Nikoletseas and J.D.P. Rolim, eds) Springer Lecture Notes in Computer Science 4240, pp. 176-183, 2006.

[10] A. Frank. Connections in combinatorial optimization, Oxford University Press, 2011.

[11] A. García and J. Tejel. Augmenting the rigidity of a graph in R^2, *Algorithmica*, February 2011, Volume 59, Issue 2, pp 145-168.

[12] S. Gortler, A. Healy, and D. Thurston. Characterizing generic global rigidity, *American Journal of Mathematics*, Volume 132, Number 4, August 2010, pp. 897-939.

[13] B. Hendrickson. Conditions for unique graph realizations, *SIAM J. Comput.* 21, 65-84, 1992.

[14] B. Jackson. Notes on the rigidity of graphs. Lecture notes, Levico, 2007.

[15] B. Jackson and T. Jordán. Connected rigidity matroids and unique realizations of graphs. *J. Combinatorial Theory Ser B*, 94:1–29, 2005.

[16] B. Jackson and T. Jordán. Graph theoretic techniques in the analysis of uniquely localizable sensor networks, in: Localization Algorithms and Strategies for Wireless Sensor Networks (G. Mao and B. Fidan, eds.), IGI Global, 2009.

[17] B. Jackson and T. Jordán. On the rigidity of molecular graphs, *Combinatorica* 28 (6), 645-658, 2008.

[18] T. Jordán. Rigid and globally rigid graphs with pinned vertices, in: Bolyai Society Mathematical Studies, 20, G.O.H. Katona, A. Schrijver, T. Szőnyi, eds., Fete of Combinatorics and Computer Science, 2010, Springer, pp. 151-172.

[19] T. Jordán. Combinatorial rigidity: graphs and matroids in the theory of rigid frameworks. In: Discrete Geometric Analysis, *MSJ Memoirs*, vol. 34, pp. 33-112, 2016.

[20] T. Jordán and W. Whiteley. Global rigidity, *in* J. E. Goodman, J. O'Rourke, and Cs. D. Tóth (eds.), *Handbook of Discrete and Computational Geometry,* 3rd ed., CRC Press, Boca Raton, pp. 1661-1694.

[21] N. Katoh and S. Tanigawa. A proof of the molecular conjecture. *Disc. Computational Geometry*, 45, 647-700, 2011.

[22] C. Király. Rigid graphs and an augmentation problem, Tech. report 2015-03, Egerváry Research Group, Budapest, 2015.

[23] G. Kortsarz and Z. Nutov. Approximating minimum-cost connectivity problems, in: Handbook of approximation algorithms and metaheuristics, T.F. Gonzalez (ed.), CRC, 2007.

[24] G. Laman. On graphs and rigidity of plane skeletal structures, *J. Engineering Math.* 4 (1970), 331-340.

[25] L. Lovász and Y. Yemini. On generic rigidity in the plane. *SIAM J. Algebraic Discrete Methods* 3 (1982), no. 1, 91–98.

[26] C.St.J.A. Nash-Williams. Decomposition of finite graphs into forests, *The Journal of the London Mathematical Society* 39 (1964) 12.

[27] J.B. Saxe. Embeddability of weighted graphs in k-space is strongly NP-hard, Tech. Report, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA, 1979.

[28] A. Schrijver. Combinatorial optimization: polyhedra and efficiency, Algorithms and combinatorics Vol. 24, Springer, 2003.

[29] W. Whiteley. Some matroids from discrete applied geometry. In J. Bonin, J. Oxley, and B. Servatius, editors, *Matroid Theory*, volume 197 of *Contemp. Math.*, pages 171–311. Amer. Math. Soc., Providence, 1996.

[30] W. Whiteley. Rigidity of molecular structures: generic and geometric analysis. In P.M. Duxbury and M.F. Thorpe, editors, *Rigidity Theory and Applications*, Kluwer/Plenum, New York, 1999.

[31] C. Yu and B.D.O. Anderson. Development of redundant rigidity theory for formation control, *Int. J. Robust and Nonlinear Control*, Vol. 19, Issue 13, 2009, pp. 1427-1446.