Discrete Comput Geom (2015) 54:339–367
DOI 10.1007/s00454-015-9692-z

Frameworks with Forced Symmetry I: Reflections and Rotations

Justin Malestein[1] · **Louis Theran**[2]

Received: 28 October 2012 / Revised: 14 March 2015 / Accepted: 23 March 2015 /
Published online: 17 June 2015

Abstract We give a combinatorial characterization of generic frameworks that are minimally rigid under the additional constraint of maintaining symmetry with respect to a finite order rotation or a reflection. To establish these results, we develop a new technique for deriving linear representations of sparsity matroids on colored graphs and extend the direction network method of proving rigidity characterizations to handle reflections.

Keywords Combinatorial rigidity · Matroids · Forced symmetry · Colored graph

1 Introduction

A Γ-*framework* is a planar structure made of *fixed-length bars* connected by *universal joints* with full rotational freedom. Additionally, the bars and joints are symmetric with respect to the action of a point group Γ. The allowed motions preserve the *length* and *connectivity* of the bars and *symmetry* with respect to the same group Γ. This model is very similar in spirit to that of the *periodic frameworks*, introduced in [4], that have recently received a lot of attention, motivated mainly by applications to zeolites [22,25]; see [16] for a discussion of the history.

Editor in Charge: János Pach

Louis Theran
louis.theran@aalto.fi

Justin Malestein
justinmalestein@gmail.com

[1] Mathematisches Institut der Universität Bonn, Bonn, Germany

[2] Aalto Science Institute and Department of Computer Science, Aalto University, Espoo, Finland

When all the allowed motions are Euclidean isometries, a framework is *rigid* and otherwise it is *flexible*. In this paper, we give a *combinatorial* characterization of minimally rigid, generic Γ-frameworks when Γ is either a finite group of rotations or a 2-element group generated by a reflection. Thus, $\Gamma \subset \mathrm{O}(2)$ and acts linearly on $\mathbb{R}^2$ in the natural way. To minimize notation, we call these *cone frameworks* and *reflection frameworks*, respectively. Since these Γ are isomorphic to some $\mathbb{Z}/k\mathbb{Z}$ for an integer $k \geq 2$, we also make this identification from now on.

1.1 The Algebraic Setup and Combinatorial Model

Formally, a Γ-framework is given by a triple $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$, where $\tilde{G}$ is a finite graph, φ is a Γ-action on $\tilde{G}$ that is free on the vertices and edges, and $\tilde{\boldsymbol{\ell}} = (\ell_{ij})_{ij \in E(\tilde{G})}$ is a vector of positive *edge lengths* assigned to the edges of $\tilde{G}$. A *realization* $\tilde{G}(\mathbf{p})$ is an assignment of points $\mathbf{p} = (\mathbf{p}_i)_{i \in V(\tilde{G})}$ such that

$$\|\mathbf{p}_j - \mathbf{p}_i\|^2 = \ell_{ij}^2 \qquad \text{for all edges } ij \in E(\tilde{G}), \tag{1}$$

$$\mathbf{p}_{\varphi(\gamma)\cdot i} = \gamma \cdot \mathbf{p}_i \qquad \text{for all } \gamma \in \mathbb{Z}/k\mathbb{Z} \text{ and } i \in V(\tilde{G}). \tag{2}$$

The set of all realizations is defined to be the *realization space* $\mathcal{R}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$. Since we require symmetry with respect to a fixed subgroup $\Gamma \subset \mathrm{Euc}(2)$ and not all rigid motions of the points preserve the symmetry, the correct notion of configuration space is $\mathcal{C}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}}) = \mathcal{R}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})/\operatorname{Cent}(\Gamma)$, where $\operatorname{Cent}(\Gamma)$ is the centralizer of Γ in $\mathrm{Euc}(2)$. A realization is *rigid* if it is isolated in the configuration space and otherwise *flexible*. A realization is *minimally rigid* if it is rigid but ceases to be so after removing any Γ-orbit of edges from $\tilde{G}$. This definition of rigidity corresponds to the intuitive one given above, since a result of Milnor [20, Lem. 3.1] implies that if a point is not isolated in the configuration space, there is a smooth path through it. An equivalent formulation of rigidity is that the only continuous length-preserving, symmetry-preserving motions are "trivial". In this case, the trivial motions are precisely $\operatorname{Cent}(\Gamma)$ which is one-dimensional for both rotational and reflective symmetries.

As the combinatorial model for cone and reflection frameworks, it will be more convenient to use colored graphs. Similarly [2,16,33], we define a *colored graph*[1] $(G, \boldsymbol{\gamma})$ to be a finite, directed graph G, with an assignment $\boldsymbol{\gamma} = (\gamma_{ij})_{ij \in E(G)}$ of an element of a group Γ to each edge. A straightforward specialization of covering space theory (see, e.g., [14, Sect. 9]) associates $(\tilde{G}, \varphi)$ with a colored graph $(G, \boldsymbol{\gamma})$: G is the quotient of $\tilde{G}$ by Γ, and the colors encode the covering map $\tilde{G} \to G$ via a natural map $\rho : \pi_1(G, b) \to \Gamma$. In this setting, the choice of base vertex does not matter, and indeed, we may define $\rho : \mathrm{H}_1(G, \mathbb{Z}) \to \mathbb{Z}/k\mathbb{Z}$ and obtain the same theory. See Sect. 2.1 for the definition of ρ via the colors.

Remark For our main theorems, we require that the symmetry group acts freely on vertices and edges. Removing this restriction, one can realize, in the case of rotational

[1] Colored graphs in this sense are also called "gain graphs" in the literature, e.g. [36].

symmetry, a framework where the group acts with a fixed vertex and with inverted edges. In this case, it is easy to reduce problems of rigidity for non-free actions to free actions, and we do so in Sect. 4.3. In the case of reflection symmetry, we can similarly extend our rigidity results to actions with inverted edges (but not fixed vertices or edges) by reducing to free actions. See the last remark in Sect. 6.2.

1.2 Main Theorems

We can now state the main results of this paper. The *cone-Laman* and *reflection-Laman graphs* appearing in the statement are defined in Sect. 2.

Theorem 1 *A generic rotation framework is minimally rigid if and only if its associated colored graph is cone-Laman.*

Theorem 2 *A generic reflection framework is minimally rigid if and only if its associated colored graph is reflection-Laman.*

Genericity has its standard meaning from algebraic geometry: the set of non-generic realizations $G(\mathbf{p})$ is a proper algebraic subset of the potential choices for $\mathbf{p}$. Whether a graph is cone-Laman and reflection-Laman can be checked by efficient combinatorial algorithms [1], making Theorems 1 and 2 good characterizations as well as entirely analogous to the Maxwell–Laman Theorem [12,19], which gives the combinatorial characterization for generic bar-joint frameworks in the plane.

Remark The approach to genericity is the one from [31]. We choose it because it is more computationally effective than ones based on algebraic independence and is also less technical. Since genericity is a Zariski-open condition, standard results imply that the non-generic set of $\mathbf{p}$ has measure zero; for a generic $\mathbf{p}$, there is an open neighborhood $U \ni \mathbf{p}$ that contains only generic points; every generic $\mathbf{p}$ is a regular point of the rigidity map that sends $\mathbf{p}$ to the edge lengths in $G(\mathbf{p})$ (or, indeed, any other polynomial map); the rank of the linear system for infinitesimal motions (described below) is maximal over all choices of $\mathbf{p}$. Some of these are sometimes taken to be defining properties of the generic set in the literature. The relevant algebraic geometry facts are collected in [11, Appendix A].

1.3 Infinitesimal Rigidity and Direction Networks

As in all known proofs of "Maxwell-Laman-type" theorems such as Theorems 1 and 2, we give a combinatorial characterization of a linearization of the problem known as *generic infinitesimal rigidity*. Given a realization $\tilde{G}(\mathbf{p})$, a set of vectors $\mathbf{v} = (\mathbf{v}_i)_{i \in V(\tilde{G})}$ is an *infinitesimal motion* of $\tilde{G}(\mathbf{p})$ if it satisfies

$$\langle \mathbf{p}_j - \mathbf{p}_i, \mathbf{v}_j - \mathbf{v}_i \rangle = 0 \qquad \text{for all edges } ij \in E(\tilde{G}). \tag{3}$$

$$\mathbf{v}_{\varphi(\gamma)\cdot i} = \gamma \cdot \mathbf{v}_i \qquad \text{for all } \gamma \in \mathbb{Z}/k\mathbb{Z} \tag{4}$$

The system (3)–(4) arises by computing the formal differential of (1)–(2). Geometrically, infinitesimal motions preserve the edge lengths to first order and preserve the symmetry of $\tilde{G}(\mathbf{p})$. A cone or reflection framework is *infinitesimally rigid* if the space of infinitesimal motions is precisely that induced by the trivial motions, i.e. $\mathrm{Cent}(\Gamma)$. Equivalently, the framework is infinitesimally rigid if and only if the space of infinitesimal motions is 1-dimensional. As discussed above, the rank of the system (3)–(4), which defines a rigidity matrix, is maximal at any generic point, so infinitesimal rigidity and rigidity coincide generically. In Sect. 4.1, we give a description of the non-generic set for cone frameworks; the case of reflection frameworks is analogous.

To characterize infinitesimal rigidity, we use a *direction network* method (cf. [16, 31,34]). A Γ-*direction network* $(\tilde{G}, \varphi, \mathbf{d})$ is a graph $\tilde{G}$ equipped with a free action $\varphi : \Gamma \to \mathrm{Aut}(G)$ and a symmetric assignment of directions with one direction $\mathbf{d}_{ij}$ per edge. The *realization space* of a direction network is the set of solutions $\tilde{G}(\mathbf{p})$ to the system of equations:

$$\left\langle \mathbf{p}_j - \mathbf{p}_i, \mathbf{d}_{ij}^{\perp} \right\rangle = 0 \qquad \text{for all edges } ij \in E(\tilde{G}) \tag{5}$$

$$\mathbf{p}_{\varphi(\gamma)\cdot i} = \gamma \cdot \mathbf{p}_i \qquad \text{for all } \gamma \in \mathbb{Z}/k\mathbb{Z} \text{ and } i \in V(\tilde{G}) \tag{6}$$

where, as above, Γ is a finite group of rotations about the origin or the 2-element group generated by a reflection. We call such a direction network a *cone direction network* in the former case and a *reflection direction network* in the latter case. By (6), a Γ-direction network is determined by assigning a direction to each edge of the colored quotient graph $(G, \boldsymbol{\gamma})$ of $(\tilde{G}, \varphi)$ (cf. [14, Lem. 17.2]). A realization of a Γ-direction network is *faithful* if none of the edges of its graph have coincident endpoints and *collapsed* if all the endpoints coincide.

A basic fact in the theory of finite planar frameworks [7,31,34] is that, if a direction network has faithful realizations, the dimension of the realization space is equal to that of the space of infinitesimal motions of a generic framework with the same underlying graph. This is also true for Γ-frameworks, if the symmetry group contains only orientation-preserving elements, as in [16], or the finite order rotations considered here. Thus, a characterization of generic cone direction networks with a 1-dimensional space of faithful realizations implies a characterization of generic minimal rigidity by a straightforward sequence of steps. We will show:

Theorem 3 *A generic cone direction network has a faithful realization that is unique up to scaling if and only if its associated colored graph is cone-Laman.*

From this, Theorem 1 follows using a slight modification of the arguments in [16, Sect. 17–18], which is presented in Sect. 4 where we will highlight where the proof breaks down for the reflection case. The main novelty in the proof of Theorem 3 is that we make a direct geometric argument for the key Proposition 3.1, since standard results on linear representations of matroid unions do not apply to the system (5)–(6).

The situation for reflection direction networks is more complicated. The reasoning used to transition from direction networks to infinitesimal rigidity in the orientation-preserving case does *not* apply verbatim in the presence of reflections. Thus, we will need to rely on a more technical analogue of Theorem 3, which we state after giving an important definition.

Let $(\tilde{G}, \varphi, \mathbf{d})$ be a direction network and define $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ to be the direction network with $(\mathbf{d}^{\perp})_{ij} = (\mathbf{d}_{ij}^{\perp})$. These two direction networks form a *special pair* if:

- $(\tilde{G}, \varphi, \mathbf{d})$ has a, unique up to scale and (vertical) translation, faithful realization.
- $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ has only collapsed realizations.

Theorem 4 *Let $(G, \boldsymbol{\gamma})$ be a colored graph with n vertices, $2n - 1$ edges, and lift $(\tilde{G}, \varphi)$. Then there are directions $\mathbf{d}$ such that the direction networks $(\tilde{G}, \varphi, \mathbf{d})$ and $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ are a special pair if and only if $(G, \boldsymbol{\gamma})$ is reflection-Laman.*

Briefly, we will use Theorem 4 as follows: the faithful realization of $(\tilde{G}, \varphi, \mathbf{d})$ gives a symmetric immersion of the graph $\tilde{G}$ that can be interpreted as a framework, and the fact that $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ has only collapsed realizations will imply that the only symmetric infinitesimal motions of this framework correspond to translation parallel to the reflection axis.

1.4 Related Work

This paper is part of a sequence extending our results about periodic frameworks [16], and the results (and proofs) reported here have been previously circulated in other preprints: Theorems 1 and 3 in [14], and Theorems 2 and 4 in [15].

This paper deals with the setting of "forced symmetry" in which all the motions considered preserve the Γ-symmetry of the framework. Some directly related results are those of Jordán et al. [10], who substantially generalize the results here to dihedral groups of order $2(2k + 1)$ using mainly inductive constructions. Tanigawa [32] proves results for scene analysis and body-bar frameworks symmetric with respect to a large number of groups using a more matroidal method.

Much of the interest in forced symmetry arises from the study of periodic frameworks, which are symmetric with respect to a translation lattice. These appear in [34], though recent interest should be traced to the more general setup in [4]. A Maxwell–Laman-type theorem for periodic frameworks in dimension 2 appears in our paper [16]. Periodic frameworks admit a number of natural variants; in dimension 2 there are combinatorial rigidity characterizations for the fixed lattice [23], partially fixed lattice with 1-degree of freedom [21] and fixed-area unit cell [17]. This paper emerged from a project to extend the rigidity characterizations from [16] to more symmetry groups; the sequel [18] has results for the orientation-preserving plane groups.

In dimensions $d \geq 3$, combinatorial rigidity characterizations for even finite bar-joint frameworks are not known. The forced-symmetric case does not appear to be much easier. However, Maxwell-type counting heuristics have been determined for a large number of space and point groups by Ross et al. [24].

All the combinatorial work mentioned so far approaches forced-symmetric frameworks in a similar formalism. Another approach used in [3,5] is to study what, in our terminology, is the underlying graph of the colored graph. This leads to a theory of a somewhat different flavor.

Another direction in the study of symmetric frameworks is to not force the symmetry constraint. This is the approach taken by Fowler and Guest [9], and a number of combinatorial characterizations are known [26–29].

1.5 Notations and Terminology

In this paper, all graphs $G = (V, E)$ may be multi-graphs. Typically, the number of vertices, edges, and connected components are denoted by n, m, and c, respectively. The notation for a colored graph is $(G, \boldsymbol{\gamma})$, and a symmetric graph with a free $\mathbb{Z}/k\mathbb{Z}$-action is denoted by $(\tilde{G}, \varphi)$. If $(\tilde{G}, \varphi)$ is the lift of $(G, \boldsymbol{\gamma})$, we denote the fiber over a vertex $i \in V(G)$ by $\tilde{i}_\gamma$, with γ ranging over $\mathbb{Z}/k\mathbb{Z}$. The fiber over a directed edge $ij \in E(G)$ with color γ_{ij} consists of the edges $\tilde{i}_\gamma \tilde{j}_{\gamma+\gamma_{ij}}$ for γ ranging over $\mathbb{Z}/k\mathbb{Z}$.

We also use (k, ℓ)-*sparse graphs* [13] and their generalizations. For a graph G, a (k, ℓ)-*basis* is a maximal (k, ℓ)-sparse subgraph; a (k, ℓ)-*circuit* is an edge-wise minimal subgraph that is not (k, ℓ)-sparse; and a (k, ℓ)-*component* is a maximal subgraph that has a spanning (k, ℓ)-graph. (To simplify terminology, we follow a convention of [31] and refer to (k, ℓ)-tight graphs simply as (k, ℓ)-graphs.)

Points in $\mathbb{R}^2$ are denoted by $\mathbf{p}_i = (x_i, y_i)$, indexed sets of points by $\mathbf{p} = (\mathbf{p}_i)$, and direction vectors by $\mathbf{d}$ and $\mathbf{v}$. For any vector $\mathbf{v}$, we denote its counter-clockwise rotation by $\pi/2$ by $\mathbf{v}^\perp$. Realizations of a cone or reflection direction network $(\tilde{G}, \varphi, \mathbf{d})$ are written as $\tilde{G}(\mathbf{p})$, as are realizations of abstract Γ-frameworks. The type of realization under consideration will always be clear from context.

2 Cone- and Reflection-Laman Graphs

In this section we define cone-Laman and reflection-Laman graphs and develop the properties we need. We start by recalling some general facts about colored graphs and the associated map ρ.

2.1 The Map ρ and Equivalent Colorings

Let $(G, \boldsymbol{\gamma})$ be a $\mathbb{Z}/k\mathbb{Z}$-colored graph. Since all the colored graphs in this paper have $\mathbb{Z}/k\mathbb{Z}$ colors, from now on we make this assumption and write simply "colored graph".

The map $\rho : \mathrm{H}_1(G, \mathbb{Z}) \to \mathbb{Z}/k\mathbb{Z}$ is defined on oriented cycles by adding up the colors on the edges; if the cycle traverses the edge in reverse, then the edge's color is added with a minus sign; otherwise it is added without a sign. As the notation suggests, ρ extends to a homomorphism from $\mathrm{H}_1(G, \mathbb{Z})$ to $\mathbb{Z}/k\mathbb{Z}$, and it is well defined even if G is not connected [16, Sect. 2]. The ρ-image of a colored graph is defined to be *trivial* if it contains only the identity.

We say $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ are *equivalent colorings* of a graph G if the corresponding lifts are isomorphic covers of G. Equivalently, $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ are equivalent if the induced representations $\mathrm{H}_1(G, \mathbb{Z}) \to \mathbb{Z}/k\mathbb{Z}$ are identical. Since all of the characteristics of rigidity or direction networks for a colored graph can be defined on the lift, $(G, \boldsymbol{\gamma})$ and $(G, \boldsymbol{\eta})$ have the same generic rigidity properties and the same generic rank for direction networks. We record the following lemma which will simplify the exposition of some later proofs.

Lemma 2.1 *Suppose G' is a subgraph of $(G, \boldsymbol{\gamma})$ such that $\rho(\mathrm{H}_1(G', \mathbb{Z}))$ is trivial. Then, there is an equivalent coloring $\boldsymbol{\eta}$ such that $\eta_{ij} = 0$ for all edges ij in G'.*

Proof It suffices to construct colors η_{ij} such that the induced representation $\mathrm{H}_1(G, \mathbb{Z}) \to \mathbb{Z}/k\mathbb{Z}$ is equal to ρ. There is no loss of generality in assuming that G is connected. Choose a spanning forest T' of G'. Extend T' to a spanning tree T of G, and set $\eta_{ij} = 0$ for all $ij \in E(T)$. Each edge ij in $G - T$ creates a fundamental (oriented) cycle C_{ij} with respect to T. Define $\eta_{ij} = \pm\rho(C_{ij})$ with the sign being positive if the orientations of ij and C_{ij} agree and negative otherwise. Since, by hypothesis, $\rho(C) = 0$ for all cycles in G', all its edges are colored 0 by $\boldsymbol{\eta}$. □

2.2 Cone- and Reflection-Laman Graphs

Let $(G, \boldsymbol{\gamma})$ be a colored graph with n vertices and m edges. We define $(G, \boldsymbol{\gamma})$ to be a *cone-Laman graph* or *reflection-Laman graph* if: $m = 2n - 1$, and for all subgraphs G', spanning n' vertices, m' edges, c' connected components with non-trivial ρ-image and c'_0 connected components with trivial ρ-image

$$m' \leq 2n' - c' - 3c'_0. \tag{7}$$

The underlying graph G of a cone-Laman graph is easily seen to be a (2, 1)-graph. In the setting of reflection symmetry, the only colored graphs arising are $\mathbb{Z}/2\mathbb{Z}$-colored; for such colorings, the families of cone-Laman and reflection-Laman graphs, which are defined purely combinatorially, are precisely the same. See Fig. 1 for some examples of various kinds of graphs described here and below.

We call $(G, \boldsymbol{\gamma})$ *cone-Laman-sparse* (resp. *reflection-Laman-sparse*) if it satisfies (7) for all subgraphs. Note that while Theorems 1 and 2 imply that these classes of graphs are matroids, at this point we have not established this. Hence, we define $(G, \boldsymbol{\gamma})$ to be a *cone-Laman-circuit* (resp. *reflection-Laman-circuit*) if it is cone-Laman-sparse (resp. reflection-Laman-sparse) after the removal of any edge.

2.3 Ross Graphs and Circuits

Another family we need is that of *Ross graphs*.[2] These are colored graphs with n vertices, $m = 2n - 2$ edges, satisfying the sparsity counts

$$m' \leq 2n' - 2c' - 3c'_0 \tag{8}$$

using the same notations as in (7). In particular, Ross graphs $(G, \boldsymbol{\gamma})$ have as their underlying graph a (2, 2)-graph G and are thus connected [13].

A *Ross circuit*[3] is a colored graph that becomes a Ross graph after removing *any* edge. The underlying graph G of a Ross-circuit $(G, \boldsymbol{\gamma})$ is a (2, 2)-circuit, and these are also known to be connected [13], so, in particular, a Ross circuit has $c'_0 = 0$ and

2 This terminology is from [1]. Elissa Ross introduced this class in [23], and we introduced this terminology in light of her contribution. In [23], they are called "constructive periodic orbit graphs".

3 The matroid of Ross graphs has more circuits, but these are the ones we are interested in here. See Sect. 2.5.

$c' = 1$, and thus satisfies (7) on the whole graph. Since (8) implies (7) and (8) holds on proper subgraphs, we see that every Ross circuit is reflection-Laman. Because reflection-Laman graphs are $(2, 1)$-graphs and subgraphs that are $(2, 2)$-sparse satisfy (8), we get the following structural result.

Proposition 2.2 ([1, Lem. 11]) *Let $(G, \boldsymbol{\gamma})$ be a reflection-Laman graph. Then each $(2, 2)$-component of G contains at most one Ross circuit, and in particular, the Ross circuits in $(G, \boldsymbol{\gamma})$ are vertex disjoint.*

2.4 Cone-(2, 2) Graphs

A colored graph $(G, \boldsymbol{\gamma})$ is defined to be a *cone*-$(2, 2)$ graph, if it has n vertices, $m = 2n$ edges, and satisfies the sparsity counts

$$m' \le 2n' - 2c'_0 \tag{9}$$

using the same notations as in (7). The link with cone-Laman graphs is the following straightforward proposition:

Proposition 2.3 *A colored graph $(G, \boldsymbol{\gamma})$ is cone-Laman if and only if, after adding a copy of any colored edge, the resulting colored graph is a cone-$(2, 2)$ graph.*

2.5 Reflection-(2, 2) Graphs

A colored graph $(G, \boldsymbol{\gamma})$ is defined to be a *reflection*-$(2, 2)$ graph, if it has n vertices, $m = 2n - 1$ edges, and satisfies the sparsity counts

$$m' \le 2n' - c' - 2c'_0 \tag{10}$$

using the same notations as in (7). The relationship between Ross graphs and reflection-$(2, 2)$ graphs we will need is:

Proposition 2.4 *Let $(G, \boldsymbol{\gamma})$ be a Ross graph. Then for either*

- *an edge ij with any color where $i \neq j$*
- *or a self-loop ℓ at any vertex i colored by* 1

the graph $(G + ij, \boldsymbol{\gamma})$ or $(G + \ell, \boldsymbol{\gamma})$ is reflection-$(2, 2)$.

Proof Adding ij with any color to a Ross $(G, \boldsymbol{\gamma})$ creates either a Ross circuit, for which $c'_0 = 0$, or a Laman-circuit with trivial ρ-image. Both of these types of graph meet this count, and so the whole of $(G + ij, \boldsymbol{\gamma})$ does as well. □

It is easy to see that every reflection-Laman graph is a reflection-$(2, 2)$ graph. The converse is not true.

Lemma 2.5 *A colored graph $(G, \boldsymbol{\gamma})$ is a reflection-Laman graph if and only if it is a reflection-$(2, 2)$ graph and no subgraph with trivial ρ-image is a $(2, 2)$-block.*

Let $(G, \boldsymbol{\gamma})$ be a reflection-Laman graph, and let $G_1, G_2, \ldots, G_t$ be the Ross circuits in $(G, \boldsymbol{\gamma})$. Define the *reduced graph* $(G^*, \boldsymbol{\gamma})$ of $(G, \boldsymbol{\gamma})$ to be the colored graph obtained by contracting each G_i, which is not already a single vertex with a self-loop (this is necessarily colored 1), into a new vertex v_i, removing any self-loops created in the process, and then adding a new self-loop with color 1 to each of the v_i. By Proposition 2.2, the reduced graph is well defined.

Proposition 2.6 *Let $(G, \boldsymbol{\gamma})$ be a reflection-Laman graph. Then its reduced graph is a reflection-$(2, 2)$ graph.*

Proof Let $(G, \boldsymbol{\gamma})$ be a reflection-Laman graph with t Ross circuits with vertex sets $V_1, \ldots, V_t$. By Proposition 2.2, the V_i are all disjoint. Now select a Ross-basis $(G', \boldsymbol{\gamma})$ of $(G, \boldsymbol{\gamma})$. The graph G' is also a $(2, 2)$-basis of G, with $2n - 1 - t$ edges, and each of the V_i spans a $(2, 2)$-block in G'. The (k, ℓ)-sparse graph Structure Theorem [13, Thm 5] implies that contracting each of the V_i into a new vertex v_i and discarding any self-loops created yields a $(2, 2)$-sparse graph G^+ on n^+ vertices and $2n^+ - 1 - t$ edges. It is then easy to check that adding a self-loop colored 1 at each of the v_i produces a colored graph satisfying the reflection-$(2, 2)$ counts (10) with exactly $2n^+ - 1$ edges. Since this is the reduced graph, we are done. □

2.6 Decomposition Characterizations

A *map-graph* is a graph with exactly one cycle per connected component. A *cone*-$(1, 1)$ or *reflection*-$(1, 1)$ graph is defined to be a colored graph $(G, \boldsymbol{\gamma})$ where G, taken as an undirected graph, is a map-graph and the ρ-image of each connected component is non-trivial. Note that by [36, Matroid Theorem], cone-$(1, 1)$ and reflection-$(1, 1)$ graphs each are bases of a matroid.

Lemma 2.7 *Let $(G, \boldsymbol{\gamma})$ be a colored graph. Then $(G, \boldsymbol{\gamma})$ is a reflection-$(2, 2)$ graph if and only if it is the union of a spanning tree and a reflection-$(1, 1)$ graph.*

Proof By [36, Matroid Theorem], reflection-$(1, 1)$ graphs are equivalent to graphs satisfying $m = n$ and

$$m' \leq n' - c'_0 \tag{11}$$

for every subgraph G'. Moreover, the right-hand side of (11) is the rank function of the matroid. We can rewrite (10) as

$$m' \leq (n' - c'_0) + (n' - c' - c'_0). \tag{12}$$

The second term in (12) is well known to be the rank function of the graphic matroid, and the lemma follows from the Edmonds-Rota construction [8] and the Matroid Union Theorem. □

Nearly the same proof yields the analogous statement for cone-$(2, 2)$ graphs.

Proposition 2.8 *A colored graph $(G, \boldsymbol{\gamma})$ is cone-$(2, 2)$ if and only if it is the union of two cone-$(1, 1)$ graphs.*

In the sequel, it will be convenient to use this slight refinement of Lemma 2.7.

Proposition 2.9 *Let* $(G, \boldsymbol{\gamma})$ *be a reflection-*$(2, 2)$ *graph. Then there is an equivalent coloring* $\boldsymbol{\gamma}'$ *of the edges of* G *such that the tree in the decomposition as in Lemma 2.7 has all edges colored by the identity.*

Proof Apply Lemma 2.1 to the tree in the decomposition. □

Proposition 2.9 has the following re-interpretation in terms of the symmetric lift $(\tilde{G}, \varphi)$:

Proposition 2.10 *Let* $(G, \boldsymbol{\gamma})$ *be a reflection-*$(2, 2)$ *graph. Then for a decomposition, as provided by Proposition 2.9, into a spanning tree* T *and a reflection-*$(1, 1)$ *graph* X*:*

- *Every edge* $ij \in T$ *lifts to the two edges* $\tilde{i}_0\tilde{j}_0$ *and* $\tilde{i}_1\tilde{j}_1$*. In particular, the vertex representatives in the lift all lie in a single connected component of the lift of* T*.*
- *Each connected component of* X *lifts to a connected graph.*

2.7 The Overlap Graph of a Cone-(2, 2) Graph

Let $(G, \boldsymbol{\gamma})$ be cone-$(2, 2)$ and fix a decomposition of it into two cone-$(1, 1)$ graphs X and Y. Let X_i and Y_i be the connected components of X and Y, respectively. Also select a base vertex x_i and y_i for each connected component of X and Y, with all base vertices on the cycle of their component. Denote the collection of base vertices by B. We define the *overlap graph* of (G, X, Y, B) to be the directed graph with:

- Vertex set B.
- A directed edge from x_i to y_j if y_j is a vertex in X_i.
- A directed edge from y_j to x_i if x_i is a vertex in Y_j

The property of the overlap graph we need is:

Proposition 2.11 *Let* $(G, \boldsymbol{\gamma})$ *be a cone-*$(2, 2)$ *graph. Fix a decomposition into cone-*$(1, 1)$ *graphs* X *and* Y *and a choice of base vertices. Then the overlap graph of* (G, X, Y, B) *has a directed cycle in each connected component.*

Proof Every vertex has exactly one incoming edge, since each vertex is in exactly one connected component of each of X and Y. Directed graphs with an in-degree exactly one have exactly one directed cycle per connected component (see, e.g., [30]). □

3 Cone Direction Networks

In this section, we prove Theorem 3. The main step in the proof is:

Proposition 3.1 *Let* $(G, \boldsymbol{\gamma})$ *be a cone-*$(2, 2)$ *graph. Then every realization of a generic direction network on* $(G, \boldsymbol{\gamma})$ *is collapsed, with all vertices placed on the rotation center.*

Next, we introduce *colored direction networks*, which will be convenient to work with.

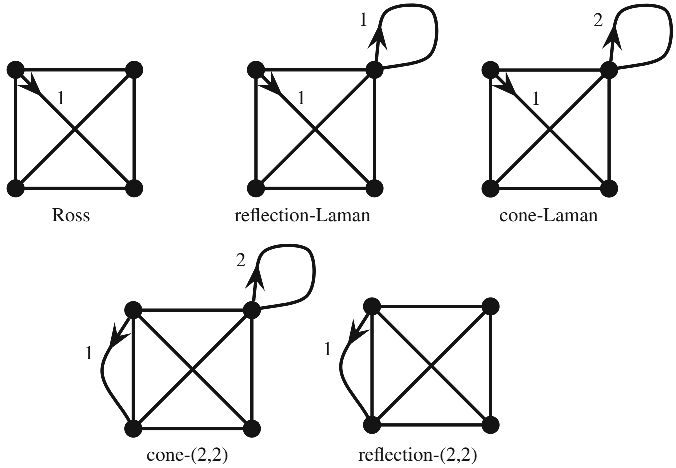

Fig. 1 Various examples of graphs. For the reflection-Laman and reflection-(2,2) graphs, the colors lie in $\mathbb{Z}/2\mathbb{Z}$. For the rest, the colors lie in $\mathbb{Z}/k\mathbb{Z}$

3.1 The Colored Realization System

The system of equations (5)–(6) defining the realization space of a cone direction network $(\tilde{G}, \varphi, \mathbf{d})$ is linear and as such has a well-defined dimension. Let $(G, \boldsymbol{\gamma})$ be the colored quotient graph of $(\tilde{G}, \varphi)$.

Since our setting is on symmetric frameworks, we require that the assigned directions also be symmetric. In other words, if $\mathbf{d}_{ij}$ is the direction for the edge $\tilde{i}_0\tilde{j}_{\gamma_{ij}}$ in the fiber of ij, then $R_k^{\gamma}\mathbf{d}_{ij}$ is the direction for the edge $\tilde{i}_\gamma\tilde{j}_{\gamma+\gamma_{ij}}$ where here and throughout this section R_k is the counter-clockwise rotation about the origin through angle $2\pi/k$. Thus, to specify a direction network on $(\tilde{G}, \varphi)$, we need only to assign a direction to one edge in each edge orbit. Furthermore, since the directions and realizations must be symmetric, the system can be reduced to the following one where the unknowns consist of n points $\mathbf{p}_i$, one for each vertex i of the quotient graph where $\mathbf{d}_{ij} = \mathbf{d}_{\tilde{i}_0\tilde{j}_{\gamma_{ij}}}$:

$$\left\langle R_k^{\gamma_{ij}} \cdot \mathbf{p}_j - \mathbf{p}_i, \mathbf{d}_{ij} \right\rangle = 0 \qquad \text{for all edges } ij \in E(G). \tag{13}$$

Proposition 3.1 can be reinterpreted to say that the colored realization system (13), in matrix form, is a linear representation for the matroid of cone-(2, 2) graphs. The representation obtained via Proposition 3.1 is different from the one produced by the matroid union construction for linearly representable matroids (see [6, Prop. 7.16.4]).

3.2 Genericity

Let $(G, \boldsymbol{\gamma})$ be a colored graph with m edges. A statement about direction networks $(\tilde{G}, \varphi, \mathbf{d})$ is *generic* if it holds on the complement of a proper algebraic subset of the possible direction assignments, which is canonically identified with $\mathbb{R}^{2m}$. Some facts about generic statements that we use frequently are as follows:

- Almost all direction assignments are generic.
- If a set of directions is generic, then so are all sufficiently small perturbations of it.
- If two properties are generic, then their intersection is as well.
- The maximum rank of (13) is a generic property.

The next proof is relatively standard.

3.3 Proof that Proposition 3.1 implies Theorem 3

We prove each direction of the statement in turn. Since it is technically easier, we prove the equivalent statement on colored direction networks.

Cone-Laman graphs generically have faithful realizations The proof in [16, Sect. 15.3] applies with small modifications.

Cone-Laman-circuits have collapsed edges For the other direction, we suppose that $(G, \boldsymbol{\gamma})$ has n vertices and is not cone-Laman. If the number of edges m is less than $2n-1$, then the realization space of any direction network is at least 2-dimensional, so it contains more than just rescalings. Thus, we assume that $m \geq 2n-1$. In this case, G is not cone-Laman- sparse, so it contains a cone-Laman-circuit $(G', \boldsymbol{\gamma})$. Thus, we are reduced to showing that any cone-Laman-circuit has, generically, only realizations with collapsed edges, since these then force collapsed edges in any realization of a generic colored direction network on $(G, \boldsymbol{\gamma})$.

There are two types of subgraphs $(G', \boldsymbol{\gamma})$ that constitute minimal violations of cone-Laman sparsity:

- $(G', \boldsymbol{\gamma})$ is a cone-(2, 2) graph.
- $(G', \boldsymbol{\gamma})$ has trivial ρ-image, and is a (2, 2)-block.

If $(G', \boldsymbol{\gamma})$ is a cone-(2, 2) graph, then Proposition 3.1 applies to it, and we are done.

For the other type, we may assume that the colors on G' are zero by Lemma 2.1 in which case the system (13) on $(G', \boldsymbol{\gamma})$ is equivalent to a direction network on G' as a finite, unsymmetric, uncolored graph. Thus, the Parallel Redrawing Theorem [35, Thm. 4.1.4] in the form [31, Thm. 3] applies directly to show that all realizations of G' have only collapsed edges.

The rest of this section proves Proposition 3.1.

3.4 Geometry of Some Generic Linear Projections

We first establish some geometric results we need below. Given a unit vector $\mathbf{v} \in \mathbb{R}^2$ and a point $\mathbf{p} \in \mathbb{R}^2$, we denote by $\ell(\mathbf{v}, \mathbf{p})$ the affine line consisting of points $\mathbf{q}$ where

$$\mathbf{q} - \mathbf{p} = \lambda \mathbf{v}$$

for some scalar $\lambda \in \mathbb{R}$, i.e. $\ell(\mathbf{v}, \mathbf{p})$ is the line through $\mathbf{p}$ in direction $\mathbf{v}$.

3.4.1 An Important Linear Equation

The following is a key lemma which will determine where certain points must lie when solving a cone direction network.

Lemma 3.2 *Suppose R is a non-trivial rotation about the origin, $\mathbf{v}^*$ is a unit vector and $\mathbf{p}$ satisfies*

$$(R - I)\mathbf{p} = \lambda \mathbf{v}^*$$

for some $\lambda \in \mathbb{R}$. Then, for some $C \in \mathbb{R}$, we have $\mathbf{p} = C\mathbf{v}$ where $\mathbf{v} = R_{\pi/2} R^{-1/2}\mathbf{v}^*$, $R^{-1/2}$ is the inverse of a square root of R, and $R_{\pi/2}$ is the counter-clockwise rotation through angle $\pi/2$.

Proof A computation shows that $(R - I)R^{-1/2} = R^{1/2} - R^{-1/2}$ is a scalar multiple of $R_{\pi/2}$. The lemma follows. □

3.4.2 The Projection $T(\mathbf{v}, \mathbf{w}, \mathbf{R})$

Let $\mathbf{v}$ and $\mathbf{w}$ be unit vectors in $\mathbb{R}^2$ and R some non-trivial rotation. Denote by $\mathbf{v}^*$ the vector $R^{1/2} \cdot \mathbf{v}^\perp$ for some choice of square root of $R^{1/2}$ of R.

We define $T(\mathbf{v}, \mathbf{w}, R)$ to be the linear projection from $\ell(\mathbf{v}, 0)$ to $\ell(\mathbf{w}, 0)$ in the direction $\mathbf{v}^*$. The following properties of $T(\mathbf{v}, \mathbf{w}, R)$ are straightforward.

Lemma 3.3 *Let $\mathbf{v}$ and $\mathbf{w}$ be unit vectors and R a non-trivial rotation. Then, the linear map $T(\mathbf{v}, \mathbf{w}, R)$:*

- *Is defined if $\mathbf{v}^*$ is not in the same direction as $\mathbf{w}$.*
- *Is identically zero if and only if $\mathbf{v}^*$ and $\mathbf{v}$ are collinear.*

3.4.3 The Scale Factor of $T(\mathbf{v}, \mathbf{w}, R)$

The image $T(\mathbf{v}, \mathbf{w}, R) \cdot \mathbf{v}$ is equal to $\lambda\mathbf{w}$ for some scalar λ. We define the *scale factor* $\lambda(\mathbf{v}, \mathbf{w}, R)$ to be this λ. We need the following elementary fact about the scaling factor of $T(\mathbf{v}, \mathbf{w}, R)$.

Lemma 3.4 *Let $\mathbf{v}$ and $\mathbf{w}$ be unit vectors such that $\mathbf{v}^*$ and $\mathbf{w}$ are linearly independent. Then the scaling factor $\lambda(\mathbf{v}, \mathbf{w}, R)$ of the linear map $T(\mathbf{v}, \mathbf{w}, R)$ is given by*

$$\frac{\langle \mathbf{v}, (\mathbf{v}^*)^\perp \rangle}{\langle \mathbf{w}, (\mathbf{v}^*)^\perp \rangle}.$$

Proof The map $T(\mathbf{v}, \mathbf{w}, R)$ is equivalent to the composition of

- perpendicular projection from $\ell(\mathbf{v}, 0)$ to $\ell((\mathbf{v}*)^\perp, 0)$, followed by
- the inverse of perpendicular projection $\ell(\mathbf{w}, 0) \to \ell((\mathbf{v}^*)^\perp, 0)$.

The first map scales the length of vectors by $\langle \mathbf{v}, (\mathbf{v}*)^\perp \rangle$ and the second by the reciprocal of $\langle \mathbf{w}, (\mathbf{v}*)^\perp \rangle$. □

From Lemma 3.4 it is immediate that

Lemma 3.5 *The scaling factor $\lambda(\mathbf{v}, \mathbf{w}, R)$ is identically 0 precisely when R is an order-2 rotation. If R is not an order-2 rotation, then $\lambda(\mathbf{v}, \mathbf{w}, R)$ approaches infinity as $\mathbf{v}^*$ approaches $\pm\mathbf{w}$.*

3.4.4 Generic Sequences of the Map $T(\mathbf{v}, \mathbf{w}, R)$

Let $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$ be unit vectors, and $S_1, S_2, \ldots, S_n$ be rotations of the form R_k^i, where R_k is a rotation of order k. We define the linear map $T(\mathbf{v}_1, S_1, \mathbf{v}_2, S_2, \ldots, \mathbf{v}_n, S_n)$ to be

$$\begin{aligned} &T(\mathbf{v}_1, S_1, \mathbf{v}_2, S_2, \ldots, \mathbf{v}_n, S_n) \\ &\quad = T(\mathbf{v}_n, \mathbf{v}_1, S_n) \circ T(\mathbf{v}_{n-1}, \mathbf{v}_n, S_{n-1}) \circ \cdots \circ T(\mathbf{v}_1, \mathbf{v}_2, S_1). \end{aligned}$$

The following proposition plays a key role in the next section, where it is interpreted as providing a genericity condition for cone direction networks.

Proposition 3.6 *Let $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$ be pairwise linearly independent unit vectors and $S_1, S_2, \ldots, S_n$ be rotations of the form R_k^i. If the $\mathbf{v}_i$ avoid a proper algebraic subset of $(\mathbb{S}^1)^n$ (that depends on the S_i), the map $T(\mathbf{v}_1, S_1, \mathbf{v}_2, S_2, \ldots, \mathbf{v}_n, S_n)$ scales the length of vectors by a factor of $\lambda \neq 1$.*

Proof If any of the $T(\mathbf{v}_i, \mathbf{v}_{i+1}, S_i)$ are identically zero, we are done, so we assume none of them are. The map $T(\mathbf{v}_1, S_1, \mathbf{v}_2, S_2, \ldots, \mathbf{v}_n, S_n)$ then scales vectors by a factor of

$$\lambda(\mathbf{v}_1, \mathbf{v}_2, S_1) \cdot \lambda(\mathbf{v}_2, \mathbf{v}_3, S_2) \cdots \lambda(\mathbf{v}_{n-1}, \mathbf{v}_n, S_{n-1}) \cdot \lambda(\mathbf{v}_n, \mathbf{v}_1, S_n)$$

which we denote by λ. That λ is constantly one is a polynomial statement in the $\mathbf{v}_i$ by Lemma 3.4, and so it is either always true or holds only on a proper algebraic subset of $\left(\mathbb{S}^1\right)^n$. This means it suffices to prove that there is one selection for the $\mathbf{v}_i$ where $\lambda \neq 1$.

Without loss of generality, we may assume that $T(\mathbf{v}_j, \mathbf{v}_{j+1}, S_j)$ is defined for all j. Note that $|\lambda(\mathbf{v}_n, \mathbf{v}_1, S_n)|$ attains an absolute (non-zero) minimum when $\mathbf{v}_n^* = \mathbf{v}_1^\perp$ by Lemma 3.4 and $|\lambda(\mathbf{v}_1, \mathbf{v}_2, S_1)|$ grows arbitrarily large as $\mathbf{v}_1^*$ moves towards $\mathbf{v}_2$. Thus $|\lambda|$ grows arbitrarily large as $\mathbf{v}_1^*$ moves towards $\mathbf{v}_2$. This means that λ is not constantly 1. □

3.5 Proof of Proposition 3.1

It will suffice to prove the existence of a single assignment of directions $\mathbf{d}$ to a cone-(2, 2) graph $(G, \boldsymbol{\gamma})$ for which we can show that all the realizations $G(\mathbf{p})$ of the resulting direction network are collapsed. The generic statement is then immediate. The strategy is to first decompose $(G, \boldsymbol{\gamma})$ into cone-(1, 1) graphs using Proposition 2.8, assign directions to each connected cone-(1, 1) graph which force a certain local geometric structure on $G(\mathbf{p})$ (Sect. 3.5.1), and then to show that the properties of the overlap graph imply that the whole realization must be collapsed (Sects. 3.5.2, 3.5.3).

3.5.1 Special Direction Networks on Connected Map-Graphs

Let $(G, \boldsymbol{\gamma})$ be a $\mathbb{Z}/k\mathbb{Z}$-colored graph that is a connected cone-$(1, 1)$ graph. We select and fix a base vertex $b \in V(G)$ that is on the (unique) cycle in G. The next lemma provides the main "gadget" that we use in the proof of Proposition 3.1.

Lemma 3.7 *Let $k \geq 2$, and let $(G, \boldsymbol{\gamma})$ be a $\mathbb{Z}/k\mathbb{Z}$-colored graph that is a connected cone-$(1, 1)$ graph with a base vertex b. Let $\gamma \in \mathbb{Z}/k\mathbb{Z}$ be the ρ-image of the cycle in G, let $\mathbf{v}$ be a unit vector, and let $\mathbf{v}^* = (R_k^{\gamma/2} \cdot \mathbf{v})^{\perp}$. We can assign directions $\mathbf{d}$ to the edges of G so that, in all realizations of the corresponding direction network $(\tilde{G}, \varphi, \mathbf{d})$ on the lift:*

- *For each $\gamma' \in \mathbb{Z}/k\mathbb{Z}$, the point $\mathbf{p}_{\tilde{b}_{\gamma'}}$ lies on the line $\ell(R_k^{\gamma'} \cdot \mathbf{v}, 0)$.*
- *The rest of the points all lie on the lines $\ell(R^{\gamma'}\mathbf{v}^*, \mathbf{p}_{\tilde{b}_{\gamma'}})$ as γ' ranges over $\mathbb{Z}/k\mathbb{Z}$.*

Proof We assign directions in the lift $\tilde{G}$ of G. We start by selecting an edge $bi \in E(G)$ that is incident on the base vertex b and in the cycle in G. Then $G - bi$ is a spanning tree T of G.

Since T is contractible, it lifts to k disjoint copies of itself in $\tilde{G}$. Let $\tilde{T}$ be the copy containing $\tilde{b}_0$. Note that $\tilde{T}$ hits the fiber over every edge in G except for bi exactly one time and the fiber over every vertex exactly one time.

Recall that we only need to assign a direction to one edge in each orbit. We first assign every edge in $\tilde{T}$ the direction $\mathbf{v}^* = (R_k^{\gamma/2} \cdot \mathbf{v})^{\perp}$. From this choice of directions already, it now follows by the connectivity of T that in any realization of the cone direction network induced on the Γ-orbit of $\tilde{T}$ any point lies on some $\ell(R_k^{\gamma'}\mathbf{v}^*, \mathbf{p}_{\tilde{b}_{\gamma'}})$.

It only remains to specify a direction for some edge in the fiber of bi. Select the edge in the fiber over bi incident on the copy of i in $\tilde{T}$. Assign this edge the direction $\mathbf{v}^*$ as well. Note that this edge is necessarily incident on the vertex $\tilde{b}_{\gamma}$ since the cycle has ρ-image γ. The choice of direction on this edge implies that

$$\mathbf{p}_{\tilde{b}_{\gamma}} - \mathbf{p}_{\tilde{b}_0} = R_k^{\gamma}\mathbf{p}_{\tilde{b}_0} - \mathbf{p}_{\tilde{b}_0} = \lambda v^*$$

for some scalar λ. It now follows from Lemma 3.2, applied to the rotation R_k^{γ}, that $\mathbf{p}_{\tilde{b}_0}$ lies on $\ell(\mathbf{v}, 0)$. □

3.5.2 Proof of Proposition 3.1 for order-2 rotations

Decompose the cone-$(2, 2)$ graph $(G, \boldsymbol{\gamma})$ into two edge-disjoint (but not necessarily connected) cone-$(1, 1)$ graphs X and Y, using Proposition 2.8. The order of the ρ-image of any cycle in either X or Y is always 2, so the construction of Lemma 3.7 implies that, by assigning the same direction $\mathbf{v}$ to every edge in X, every vertex in any realization lies on a single line through the origin in the direction of $\mathbf{v}$. Similarly for edges in Y, assign a direction $\mathbf{w}$ different than $\mathbf{v}$.

Since every vertex is at the intersection of two skew lines through the origin, the proposition is proved.

3.5.3 Proof of Proposition 3.1 for rotations of order $k \geq 3$

Using Proposition 2.8, fix a decomposition of the cone-(2, 2) graph $(G, \boldsymbol{\gamma})$ into two cone-(1, 1) graphs X and Y and choose base vertices for the connected components of X and Y. This determines an overlap graph D. Since we deal with them independently of X and Y, we define G_i to be the connected components of X and Y, and recall that these partition the edges of G. We denote the base vertex of G_i by b_i. Since we use subscripts to denote the subgraph of the base vertex, in this section we will use the notation $\gamma \cdot \tilde{b}_i$ as γ ranges over $\mathbb{Z}/k\mathbb{Z}$ for the fiber over vertex b_i.

Assigning directions For each G_i, select a unit vector $\mathbf{v}_i$ such that:

- For any i, j, we have $\mathbf{v}_i \neq R_k^{\gamma} \mathbf{v}_j$ for all $\gamma \in \mathbb{Z}/k\mathbb{Z}$.
- For all choices k_i, the vectors $\mathbf{w}_i = R_k^{k_i} \mathbf{v}_i$ are generic in the sense of Proposition 3.6.

Now, for each G_i we assign directions as prescribed by Lemma 3.7 where $\mathbf{v}_i$ is the vector input into the lemma. This is well defined, since G_i partition the edges of G. (They clearly overlap on the vertices—we will exploit this fact below—but it does not prevent us from assigning edge directions independently.)

We define the resulting colored direction network to be $(G, \boldsymbol{\gamma}, \mathbf{d})$ and the lifted cone direction network $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$. We also define, as a convenience, the rotation S_i to be R_k^{γ} where γ is the ρ-image of the unique cycle in G_i.

Local structure of realizations Let G_i and G_j be distinct connected cone-(1, 1) components and suppose that there is a directed edge $b_i b_j$ in the overlap graph D. We have the following relationship between $\mathbf{p}_{\tilde{b}_i}$ and $\mathbf{p}_{\tilde{b}_j}$ in realizations of $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$.

Lemma 3.8 *Let $\tilde{G}(\mathbf{p})$ be a realization of the cone direction network $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ defined above. Let vertices b_i and b_j in $V(G)$ be the base vertices of G_i and G_j, and suppose that $b_i b_j$ is a directed edge in the overlap graph D. Let $\gamma \cdot \tilde{b}_i$ be some vertex in the fiber of b_i. Then for some γ', we have $\mathbf{p}_{\gamma' \cdot \tilde{b}_j} = T(R_k^{\gamma'} \mathbf{v}_i, R_k^{\gamma} \mathbf{v}_j, S_i) \cdot \mathbf{p}_{\gamma \cdot \tilde{b}_i}$.*

The proof is illustrated in Fig. 2.

Proof By Lemma 3.7, the vertex $\mathbf{p}_{\gamma \cdot \tilde{b}_i}$ lies on the line $\ell(R_k^{\gamma} \mathbf{v}, 0)$, and $\mathbf{p}_{\gamma' \cdot \tilde{b}_j}$ lies on $\ell(R_k^{\gamma} \mathbf{v}^*, \mathbf{p}_{\gamma \cdot \tilde{b}_i})$ for some $\gamma' \in \mathbb{Z}/k\mathbb{Z}$ since the vertex b_j lies in the map-graph G_i. By Lemma 3.7 applied to G_j, the vertex $\mathbf{p}_{\gamma' \cdot \tilde{b}_j}$ lies on the line $\ell(R_k^{\gamma'} \mathbf{v}, 0)$. This is exactly the situation captured by the map $T(R_k^{\gamma'} \mathbf{v}_i, R_k^{\gamma} \mathbf{v}_j, S_i)$. □

Base vertices on cycles in D must be at the origin Let b_i be the base vertex in G_i that is also on a directed cycle in D. The next step in the proof is to show that all representatives in b_i must be mapped to the origin in any realization of $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$.

Lemma 3.9 *Let $\tilde{G}(\mathbf{p})$ be a realization of the cone direction network $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ defined above, and let $b_i \in V(G)$ be a base vertex that is also on a directed cycle in D (one*

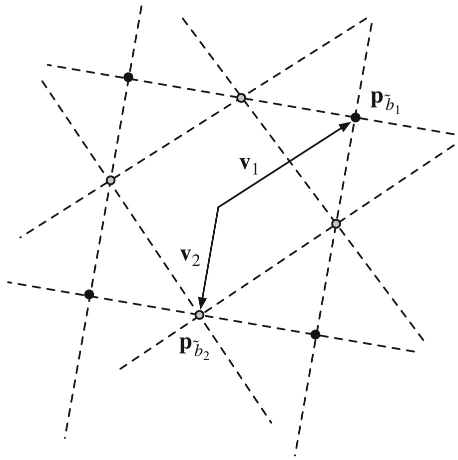

Fig. 2 Example of the local structure of the proof of Proposition 3.1; in this example, $R_4^2\mathbf{p}_{\tilde{b}_2}$ is the image of $p_{\tilde{b}_1}$ via the projection $T(\mathbf{v}_1, R_4^2\mathbf{v}_2, S_1)$

exists by Proposition 2.11). *Then all vertices in the fiber over* b_i *must be mapped to the origin.*

Proof Iterated application of Lemma 3.8 along the cycle in D containing b_i tells us that any vertex in the fiber over b_i is related to another vertex in the same fiber by a linear map meeting the hypothesis of Proposition 3.6. This implies that if any vertex in the fiber over b_i is mapped to a point not the origin, some other vertex in the same fiber would be mapped to a point at a different distance to the origin. This is a contradiction, since all realizations $\tilde{G}(\mathbf{p})$ are symmetric with respect to R_k, so in fact the fiber over b_i was mapped to the origin. □

All base vertices must be at the origin So far we have shown that every base vertex b_i that is on a directed cycle in the overlap graph D is mapped to the origin in any realization $\tilde{G}(\mathbf{p})$ of $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$. However, since every base vertex is connected to the cycle in its connected component by a directed path in D, we can show that all the base vertices are at the origin.

Lemma 3.10 *Let* $\tilde{G}(\mathbf{p})$ *be a realization of the cone direction network* $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ *defined above. Then all vertices in the fiber over* b_i *must be mapped to the origin.*

Proof The statement is already proved for base vertices on a directed cycle in Lemma 3.9. Any base vertex not on a directed cycle, say b_i, is at the end of a directed path which starts at a vertex on the directed cycle. Thus $\mathbf{p}_{\gamma \cdot b_i}$ is the image of 0 under some linear map, and hence is at the origin. □

All vertices must be at the origin The proof of Proposition 3.1 then follows from the observation that, if all the base vertices b_i must be mapped to the origin in $\tilde{G}(\mathbf{p})$, then Lemma 3.7 implies that *every* vertex in the lift of G_i lies on a family of k lines intersecting at the origin. Since every vertex is in the span of two of the G_i, and these families of lines intersect only at the origin, we are done: $\tilde{G}(\mathbf{p})$ must put all the points at the origin.

4 Infinitesimal Rigidity of Cone Frameworks

The generic rigidity of a cone framework $(\tilde{G}, \varphi, \tilde{\ell})$ is a property of the underlying colored graph. To see this, note that we can identify the realization space $\mathcal{R}(\tilde{G}, \varphi, \tilde{\ell})$ with the solutions to the following system, defined via the associated colored graph $(G, \boldsymbol{\gamma})$:

$$\|R_k^{\gamma_{ij}} \mathbf{p}_j - \mathbf{p}_i\|^2 = \ell_{ij}^2 \quad \text{for all edges } ij \in E(G). \tag{14}$$

Here ℓ_{ij} is equal to the length of any lift of the edge ij which, by symmetry, is independent of the lift. We call the resulting object a *colored framework* $(G, \boldsymbol{\gamma}, \ell)$ on the quotient graph and denote its realization space by $\mathcal{R}(G, \boldsymbol{\gamma}, \ell)$. Since the two spaces have the same dimension, for any choice $V \subset V(\tilde{G})$ of vertex–orbit representatives in $\tilde{G}$, the projection onto the set $(\mathbf{p}_i)_{i \in V}$ induces an algebraic isomorphism.

Let G have n vertices. Computing the formal differential of (14), we see that a vector $\mathbf{v} \in \mathbb{R}^{2n}$ is an infinitesimal motion if and only if

$$\langle R_k^{\gamma_{ij}} \mathbf{v}_j - \mathbf{v}_i, R_k^{\gamma_{ij}} \mathbf{p}_j - \mathbf{p}_i \rangle = 0 \quad \text{for all edges } ij \in E(G). \tag{15}$$

We define $(G, \boldsymbol{\gamma}, \ell)$ to be *infinitesimally rigid* if the system (15) has rank $2n - 1$. It is easy to see that infinitesimal rigidity of $(G, \boldsymbol{\gamma}, \ell)$ coincides with that of the lift $(\tilde{G}, \varphi, \tilde{\ell})$.

4.1 Framework Genericity

Our approach to genericity for cone frameworks is a small extension of the one for finite frameworks from [31]. Let $(G, \boldsymbol{\gamma}, \ell)$ be a colored framework. A realization $G(\mathbf{p})$ is generic if the rank of (15) is maximum among all the realizations. Thus, the set of non-generic realizations consists simply of those for which the (complexification of the) system (15) does not attain its maximal rank. This is cut out by the minors of the matrix form of (15), and so is clearly algebraic. Since the natural homeomorphism $\mathcal{R}(\tilde{G}, \varphi, \tilde{\ell}) \to \mathcal{R}(G, \boldsymbol{\gamma}, \ell)$ is an algebraic map, the pre-image in $\mathcal{R}(\tilde{G}, \varphi, \tilde{\ell})$ of the non-generic subset of $\mathcal{R}(G, \boldsymbol{\gamma}, \ell)$ is also algebraic.

4.2 Proof of Theorem 1

We prove necessity by inspecting (15) and then sufficiency with Theorem 3.

The Maxwell direction Suppose that $(G, \boldsymbol{\gamma})$ is a colored graph on $m = 2n - 1$ edges that is not cone-Laman-sparse. Thus G contains (at least) one of two possible types of cone-Laman-circuits which we call G': a cone-$(2, 2)$ graph or a $(2, 2)$-graph with trivial ρ-image. In the former case, the subgraph has n vertices and $2n$ edges, and since any framework has a trivial motion arising from rotation, the system (15) has a dependency.

In the latter case, by Lemma 2.1, we may assume the edges are colored by 0, in which case the system (15) is identical to the well-known rigidity matrix for finite frameworks. Since G' is not (2, 3)-sparse, the Maxwell-Laman Theorem provides a dependency in (15).

The Laman direction Theorem 3 implies that (13) has, generically, rank $2n - 1$ if $(G, \boldsymbol{\gamma})$ is cone-Laman. The next proposition says that (15) has the same rank.

Proposition 4.1 *Let $G(\mathbf{p})$ be a realization of a cone framework with colored graph $(G, \boldsymbol{\gamma})$. If $G(\mathbf{p})$ is faithful and solves the direction network $(G, \boldsymbol{\gamma}, \mathbf{d})$, then the system (15) for $G(\mathbf{p})$ has the same rank as (13) for $(G, \boldsymbol{\gamma}, \mathbf{d})$.*

Proof Let $R_{\pi/2}$ be the counter-clockwise rotation through angle $\pi/2$. If $G(\mathbf{p})$ faithfully solves $(G, \boldsymbol{\gamma}, \mathbf{d})$, then $R_k^{\gamma_{ij}}\mathbf{p}_j - \mathbf{p}_i = \alpha\mathbf{d}_{ij}$. A vector $\mathbf{q}$ is therefore an infinitesimal motion if and only if $\mathbf{q}^\perp$ is a solution to (13) by this computation:

$$\begin{aligned} \langle R_k^{\gamma_{ij}}\mathbf{q}_j - \mathbf{q}_i, \mathbf{d}_{ij}\rangle = 0 &\iff \\ \langle R_{\pi/2}(R_k^{\gamma_{ij}}\mathbf{q}_j - \mathbf{q}_i), R_{\pi/2}\mathbf{d}_{ij}\rangle = 0 &\iff \\ \langle R_k^{\gamma_{ij}}R_{\pi/2}\mathbf{q}_j - R_{\pi/2}\mathbf{q}_i, R_{\pi/2}\mathbf{d}_{ij}\rangle = 0. & \end{aligned}$$

Thus, (15) and (13) have isomorphic solution spaces and hence the same rank. □

Remark That $R_{\pi/2}$ and $R_k^{\gamma_{ij}}$ commute is critical in the computation used to prove Proposition 4.1. The corresponding argument for reflection frameworks would fail since $R_k^{\gamma_{ij}}$ would be replaced by a reflection which does not commute with $R_{\pi/2}$. In fact, as we will see, the ranks of the two systems can be different in the reflection case.

4.3 Rigidity for Non-free Actions

We briefly remark here on symmetric frameworks with fixed vertices or inverted edges. If there is an edge $\tilde{i}\tilde{j} \in \tilde{G}$ and a group element $\gamma \in \Gamma$ such that $\gamma \cdot \tilde{i}\tilde{j} = \tilde{j}\tilde{i}$, then $\tilde{i}$ and $\tilde{j}$ descend to the same vertex, say i, in the quotient. The inverted edge, in terms of rigidity, forces $\mathbf{p}_i$ and $R^\gamma \cdot \mathbf{p}_i$ to remain at a constant distance. This is the same constraint as if there were a self-loop at i with color γ in the quotient graph $(G, \boldsymbol{\gamma})$. Thus, for every inverted edge, we simply put such a self-loop in the quotient graph and appeal to Theorem 1.

If there is a fixed vertex, say $\tilde{i}$, then in any realization it lies at the origin. (We may assume without loss of generality that there is only one fixed vertex.) Suppose $\tilde{j}$ and consequently its Γ-orbit are connected to $\tilde{i}$ by some edge (orbit). As a symmetric framework, this forces $\mathbf{p}_j$ and its Γ-orbit to be the vertices of a regular k-gon with fixed distance from the origin. The same constraint can be enforced by deleting the edge orbit to $\tilde{i}$ and replacing it with the edges of the regular polygon (see Fig. 3). Thus, we reduce the problem again to the case of a free Γ-action.

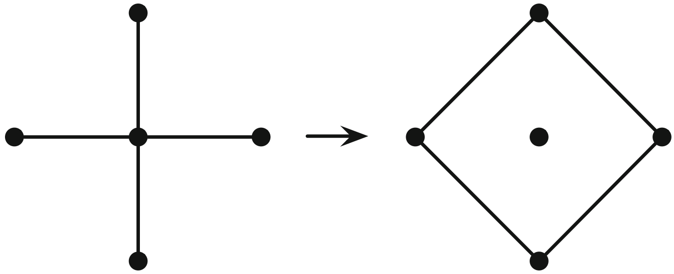

Fig. 3 Operation removing an edge orbit to a fixed vertex

5 Special Pairs of Reflection Direction Networks

We recall, from the introduction, that for reflection direction networks, $\mathbb{Z}/2\mathbb{Z}$ acts on the plane by some reflection through the origin. It is clear that we can reduce to the case where the reflection is through the y-axis, and we make this assumption for the remainder of this section.

5.1 Direction Networks on Ross Graphs

We first characterize the colored graphs for which generic direction networks have strongly faithful realizations. A realization is *strongly faithful* if no two vertices lie on top of each other. This is a stronger condition than simply being faithful which only requires that edges are not to be collapsed.

Proposition 5.1 *A generic direction network* $(\tilde{G}, \varphi, \mathbf{d})$ *has a unique, up to* (*vertical*) *translation and scaling, strongly faithful realization if and only if its associated colored graph is a Ross graph.*

To prove Proposition 5.1 we expand upon the method from [16, Sects. 17–18] and use the following proposition.

Proposition 5.2 *Let* $(G, \boldsymbol{\gamma})$ *be a reflection-*$(2, 2)$ *graph. Then a generic direction network on the symmetric lift* $(\tilde{G}, \varphi)$ *of* $(G, \boldsymbol{\gamma})$ *has only collapsed realizations.*

Since the proof of Proposition 5.2 requires a detailed construction, we first show how it implies Proposition 5.1.

5.2 Proof that Proposition 5.2 Implies Proposition 5.1

Let $(G, \boldsymbol{\gamma})$ be a Ross graph, and assign directions $\mathbf{d}$ to the edges of G such that, for any extension $(G + ij, \boldsymbol{\gamma})$ of $(G, \boldsymbol{\gamma})$ to a reflection-$(2, 2)$ graph as in Proposition 2.4, $\mathbf{d}$ can be extended to a set of directions that is generic in the sense of Proposition 5.2. This is possible because there are a finite number of such extensions.

For this choice of $\mathbf{d}$, the realization space of the direction network $(\tilde{G}, \varphi, \mathbf{d})$ is 2-dimensional. Since solutions to (13) may be scaled or translated in the vertical direction, all solutions to $(\tilde{G}, \varphi, \mathbf{d})$ are related by scaling and translation. It then follows that a pair of vertices in the fibers over i and j are either distinct from each other in all non-zero solutions to (13) or always coincide. In the latter case, adding the edge ij with any direction does not change the dimension of the solution space, no matter what direction we assign to it. It then follows that the solution spaces of generic direction networks on $(\tilde{G}, \varphi, \mathbf{d})$ and $(\widetilde{G + ij}, \varphi, \mathbf{d})$ have the same dimension, which is a contradiction by Proposition 5.2.

For the opposite direction, suppose $(G, \boldsymbol{\gamma})$ is not a Ross graph. A proof similar to that in Sect. 3.3 applies. If $m < 2n - 2$, then dimension counting tells us that the space of realizations cannot be unique up to (vertical) translations and scaling. If $m \geq 2n - 2$, then $(G, \boldsymbol{\gamma})$ has one of two types of circuits, either a reflection-(2, 2)-subgraph or a (2, 2)-subgraph with trivial ρ-image. In the former case, we are done by Proposition 5.2, and in the latter case, the same proof from Sect. 3.3 applies.

5.3 Proof of Proposition 5.2

Let $(G, \boldsymbol{\gamma})$ be a reflection-(2, 2) graph associated to $(\tilde{G}, \varphi)$. It is sufficient to construct a set of directions $\mathbf{d}$ such that the direction network $(\tilde{G}, \varphi, \mathbf{d})$ has only collapsed realizations. In the rest of the proof, we construct a set of directions $\mathbf{d}$ and then verify that the colored direction network $(G, \boldsymbol{\gamma}, \mathbf{d})$ has only collapsed solutions. The proposition then follows from the equivalence of colored direction networks with reflection direction networks.

Combinatorial decomposition We apply Proposition 2.9 to decompose $(G, \boldsymbol{\gamma})$ into a spanning tree T with all colors the identity and a reflection-(1, 1) graph X. For now, we further assume that X is connected.

Assigning directions Let $\mathbf{v}$ be a direction vector that is not horizontal or vertical. For each edge $ij \in T$, set $\mathbf{d}_{ij} = \mathbf{v}$. Assign all the edges of X the vertical direction. Denote by $\mathbf{d}$ this assignment of directions.

All realizations are collapsed We now show that the only realizations of $(\tilde{G}, \varphi, \mathbf{d})$ have all vertices on top of each other. By Proposition 2.10, T lifts to two copies of itself, in $\tilde{G}$. It then follows from the connectivity of T and the construction of $\mathbf{d}$ that, in any realization, there is a line L with direction $\mathbf{v}$ such that every vertex of $\tilde{G}$ must lie on L or its reflection. Since the vertical direction is preserved by reflection, the connectivity of the lift of X, again from Proposition 2.10, implies that every vertex of $\tilde{G}$ lies on a single vertical line, which must be the y-axis by reflection symmetry.

Thus, in any realization of $(\tilde{G}, \varphi, \mathbf{d})$ all the vertices lie at the intersection of L, the reflection of L through the y-axis and the y-axis itself. This is a single point, as desired. Figure 4 shows a schematic of this argument.

X does not need to be connected Finally, we can remove the assumption that X was connected by repeating the argument for each connected component of X separately.

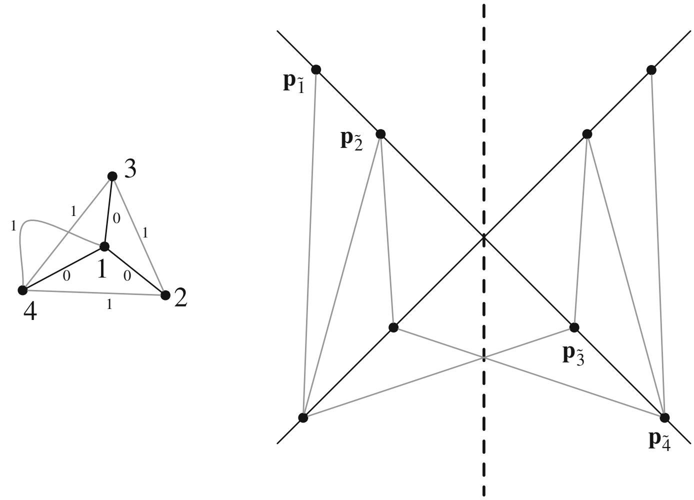

Fig. 4 Schematic for proof of Proposition 5.2 and Lemma 5.4: the y-axis is shown as a dashed line. The corresponding colored graph is depicted on the left with the tree indicated by *black lines* and the reflection-(1, 1) graph indicated by the *gray lines*. The *right-hand figure* indicates a realization where only the directions on the tree are enforced. If the *gray lines* are then forced to be vertical, the entire framework collapses to the y-axis. If the *gray lines* are forced to be horizontal, the framework takes on the form given in Fig. 5

5.4 Special Pairs for Ross Circuits

Theorem 4 will reduce to the case of a Ross circuit.

Proposition 5.3 *Let* $(G, \boldsymbol{\gamma})$ *be a Ross circuit with lift* $(\tilde{G}, \varphi)$*. Then there is an edge* $i'j'$ *with non-zero color such that for a generic direction network* $(\tilde{G}', \varphi, \mathbf{d})$ *with colored graph* $(G - i'j', \boldsymbol{\gamma})$*:*

- *For all faithful realizations of* $(\tilde{G}', \varphi, \mathbf{d})$*, we have that* $\mathbf{p}_{\tilde{j}'_1} - \mathbf{p}_{\tilde{i}'_0}$ *is a non-zero vector with direction independent of the realization. In particular,* $(\tilde{G}', \varphi, \mathbf{d})$ *induces a well-defined direction on the edge* $i'j'$*, which extends to an assignment of directions to the edges of* G*.*
- *The direction networks* $(\tilde{G}, \varphi, \mathbf{d})$ *and* $(\tilde{G}, \varphi, (\mathbf{d})^{\perp})$ *are a special pair.*

Before giving the proof, we describe the idea. We are after sets of directions that lead to faithful realizations of Ross circuits. By Proposition 5.2, these directions must be non-generic. A natural way to obtain such a set of directions is to discard an edge ij from the colored quotient graph, apply Proposition 5.1 to obtain a generic set of directions $\mathbf{d}'$ with a strongly faithful realization $\tilde{G}'(\mathbf{p})$, and then simply set the directions on the edges in the fiber over ij to be the difference vectors between the points.

Proposition 5.1 tells us that this procedure induces a well-defined direction for the edge ij, allowing us to extend $\mathbf{d}$ from G' to G in a controlled way. However, it does *not* tell us that rank of $(\tilde{G}, \varphi, \mathbf{d})$ will rise when the directions are turned by angle

$\pi/2$, and this seems hard to do directly. Instead, we construct a set of directions $\mathbf{d}$ so that $(\tilde{G}, \varphi, \mathbf{d})$ is rank deficient and has realizations where $\mathbf{p}_i \neq \mathbf{p}_j$, and $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ is generic. Then we make a perturbation argument to show the existence of a special pair.

The construction we use is, essentially, the one used in the proof of Proposition 5.2 but turned through angle $\pi/2$. The key geometric insight is that horizontal edge directions are preserved by the reflection, so the "gadget" of a line and its reflection crossing on the y-axis, as in Fig. 4, degenerates to just a single line.

5.5 Proof of Proposition 5.3

Let $(G, \boldsymbol{\gamma})$ be a Ross circuit; recall that this implies that $(G, \boldsymbol{\gamma})$ is a reflection-Laman graph.

Combinatorial decomposition We decompose $(G, \boldsymbol{\gamma})$ into a spanning tree T and a reflection-$(1, 1)$ graph X as in Proposition 2.10. In particular, we again have all edges in T colored by the identity. For now, we *assume that X is connected*, and we fix $i'j'$ to be an edge that is on the cycle in X with $\gamma_{i'j'} \neq 0$; such an edge must exist by the hypothesis that X is reflection-$(1, 1)$. Let $G' = G \setminus i'j'$. Furthermore, let $\tilde{T}_0$ and $\tilde{T}_1$ be the two connected components of the lift of T. For a vertex $i \in G$, the lift $\tilde{i}_k$ lies in $\tilde{T}_k$. We similarly denote the lifts of i' and j' by $\tilde{i}'_0, \tilde{i}'_1$ and $\tilde{j}'_0, \tilde{j}'_1$.

Assigning directions The assignment of directions is as follows: to the edges of T, we assign a direction $\mathbf{v}$ that is neither vertical nor horizontal. To the edges of X we assign the horizontal direction. Define the resulting direction network to be $(\tilde{G}, \varphi, \mathbf{d})$ and the direction network induced on the lift of G' to be $(\tilde{G}', \varphi, \mathbf{d})$.

The realization space of $(\tilde{G}, \varphi, \mathbf{d})$ Figs. 4 and 5 contains a schematic picture of the arguments that follow.

Lemma 5.4 *The realization space of $(\tilde{G}, \varphi, \mathbf{d})$ is 2-dimensional and parameterized by exactly one representative in the fiber over the vertex i selected above.*

Proof In a manner similar to the proof of Proposition 5.2, the directions on the edges of T force every vertex to lie either on a line L in the direction $\mathbf{v}$ or its reflection. Since the lift of X is connected, we further conclude that all the vertices lie on a single horizontal line. Thus, all the points $\mathbf{p}_{\tilde{j}_0}$ are at the intersection of the same horizontal line and L or its reflection. These determine the locations of the $\mathbf{p}_{\tilde{j}_1}$, so the realization space is parameterized by the location of $\mathbf{p}_{\tilde{i}'_0}$. □

Inspecting the argument more closely, we find that:

Lemma 5.5 *In any realization $\tilde{G}(\mathbf{p})$ of $(\tilde{G}, \varphi, \mathbf{d})$, all the $\mathbf{p}_{\tilde{j}_0}$ are equal and all the $\mathbf{p}_{\tilde{j}_1}$ are equal.*

Proof Because the colors on the edges of T are all zero, it lifts to two copies of itself, one of which spans the vertex set $\{\tilde{j}_0 : j \in V(G)\}$ and one which spans

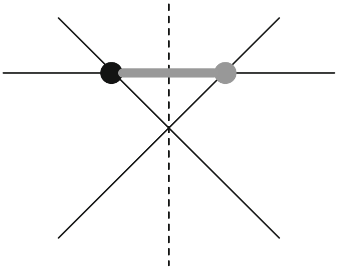

Fig. 5 Schematic of the proof of Proposition 5.3: the y-axis is shown as a *dashed line*. The directions on the edges of the lift of the tree T force all the vertices to be on one of the two lines meeting at the y-axis. The horizontal directions on the connected reflection-(1, 1) graph X force the point $\mathbf{p}_{\tilde{j}_0}$ to be at the intersection marked by the *black dot* and $\mathbf{p}_{\tilde{j}_1}$ to be at the intersection marked by the *gray* one. The *thick gray line* indicates a thick mass of horizontal edges

$\{\tilde{j}_1 : j \in V(G)\}$. It follows that in a realization, we have all the $\mathbf{p}_{\tilde{j}_0}$ on L and the $\mathbf{p}_{\tilde{j}_1}$ on the reflection of L. □

In particular, because the color $\gamma_{i'j'}$ on the edge $i'j'$ is 1, we obtain the following.

Lemma 5.6 *The realization space of* $(\tilde{G}, \varphi, \mathbf{d})$ *contains points where the fiber over the edge* $i'j'$ *is not collapsed.*

The realization space of $(\tilde{G}', \varphi, \mathbf{d})$ The conclusion of Lemma 5.4 implies that the realization system for $(\tilde{G}, \varphi, \mathbf{d})$ is rank deficient by one. Next we show that removing the edge $i'j'$ results in a direction network that has full rank on the colored graph $(G', \boldsymbol{\gamma})$.

Lemma 5.7 *The realization space of* $(\tilde{G}, \varphi, \mathbf{d})$ *is canonically identified with that of* $(\tilde{G}', \varphi, \mathbf{d})$.

Proof In the proof of Lemma 5.4, it was not essential that X lifts to a connected subgraph of $\tilde{G}$. It was only required that X spans the vertices, and this is true of $X - i'j'$. Since the two lifts of a vertex must always lie on the same horizontal line in a realization, if any lift of i and any lift of j lie on the same horizontal line, then all lifts do. It is then easy to conclude all points lie on a single horizontal line. □

The realization space of $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ Next, we consider what happens when we turn all the directions by $\pi/2$.

Lemma 5.8 *The realization space of* $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ *has only collapsed solutions.*

Proof This is exactly the construction used to prove Proposition 5.2. □

Perturbing $(\tilde{G}, \varphi, \mathbf{d})$ To summarize what we have shown so far:

(a) $(\tilde{G}, \varphi, \mathbf{d})$ has a 2-dimensional realization space parameterized by $\mathbf{p}_{\tilde{i}'_0}$ and identified with that of a full-rank direction network on the Ross graph $(G', \boldsymbol{\gamma})$.

(b) There are points $\tilde{G}(\mathbf{p})$ in this realization space where $\mathbf{p}_{\tilde{i}'_0} \neq \mathbf{p}_{\tilde{j}'_1}$.
(c) $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ has a 1-dimensional realization space containing only collapsed solutions.

What we have not shown is that the realization space of $(\tilde{G}, \varphi, \mathbf{d})$ has *faithful* realizations, since the ones we constructed all have many coincident vertices. Proposition 5.1 will imply the rest of the theorem, provided that the above properties hold for any small perturbation of $\mathbf{d}$, since some small perturbations of *any* assignment of directions to the edges of $(G', \boldsymbol{\gamma})$ have only faithful realizations.

Lemma 5.9 *Let $\hat{\mathbf{d}}'$ be a perturbation of the directions $\mathbf{d}$ on the edges of G' only. If $\hat{\mathbf{d}}'$ is sufficiently close to $\mathbf{d}|_{E(G')}$, then there are realizations of the direction network $(\tilde{G}', \varphi, \hat{\mathbf{d}}')$ such that the direction of $\mathbf{p}_{\tilde{j}'_1} - \mathbf{p}_{\tilde{i}'_0}$ is non-zero and a small perturbation of $\mathbf{d}_{ij}$.*

Proof The realization space is parameterized by $\mathbf{p}_{\tilde{i}'_0}$ (for directions sufficiently close to $\mathbf{d}'$), and so $\mathbf{p}_{\tilde{j}'_1}$ varies continuously with the directions on the edges and $\mathbf{p}_{\tilde{i}'_0}$. Since there are realizations of $(\tilde{G}', \varphi, \mathbf{d})$ with $\mathbf{p}_{\tilde{i}_0} \neq \mathbf{p}_{\tilde{j}_1}$, the lemma follows. □

Lemma 5.9 implies that any sufficiently small perturbation of the directions assigned to the edges of G' gives a direction network that induces a well-defined direction on the edge $i'j'$ which is itself a small perturbation of $\mathbf{d}_{i'j'}$. Since the ranks of $(\tilde{G}', \varphi, \mathbf{d}')$ and $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ are stable under small perturbations, this implies that we can perturb $\mathbf{d}$ to a $\hat{\mathbf{d}}$ so that $\hat{\mathbf{d}}|_{E(G')}$ is generic in the sense of Proposition 5.1, while preserving faithful realizability of $(\tilde{G}, \varphi, \hat{\mathbf{d}})$ and full rank of the realization system for $(\tilde{G}, \varphi, \hat{\mathbf{d}}^{\perp})$. The Proposition is proved when X is connected.

X need not be connected The proof is then complete once we remove the additional assumption that X was connected. Let X have connected components $X_1, X_2, \ldots, X_c$. For each of the X_i, we can identify an edge $(i'j')_k$ with the same properties as $i'j'$ above.

Assign directions to the tree T as above. For X_1, we assign directions exactly as above. For each of the X_k with $k \geq 2$, we assign the edges of $X_k \setminus (i'j')_k$ the horizontal direction and $(i'j')_k$ a direction that is a small perturbation of horizontal.

With this assignment $\mathbf{d}$, we see that for any realization of $(\tilde{G}, \varphi, \mathbf{d})$, each of the X_k, for $k \geq 2$, is realized as completely collapsed to a single point at the intersection of the line L and the y-axis. Moreover, in the direction network on $\mathbf{d}^{\perp}$, the directions on these X_i are a small perturbation of the ones used on X in the proof of Proposition 5.2. From this it follows that any realization of $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ is completely collapsed and hence full rank.

We now see that this new set of directions has properties (a), (b), and (c) above required for the perturbation argument. Since that argument makes no reference to the decomposition, it applies verbatim to the case where X is disconnected.

5.6 Proof of Theorem 4

The easier direction to check is necessity.

The Maxwell direction If $(G, \boldsymbol{\gamma})$ is not reflection-Laman, then it contains either a Laman-circuit with trivial ρ-image or a violation of (2, 1)-sparsity. A violation of (2, 1)-sparsity implies that the realization system (13) of $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ has a dependency, since the realization space is always at least 1-dimensional.

Suppose instead there is a Laman-circuit G' with trivial ρ-image. Then any direction network on $(G', \boldsymbol{\gamma})$ is equivalent to a direction network on the (finite, uncolored, non-symmetric) graph G'. (The lift $\tilde{G}'$ is two mirror images of G'.) In this case, similar to Proposition 4.1, $(G', \boldsymbol{\gamma}, \mathbf{d})$ and $(G', \boldsymbol{\gamma}, \mathbf{d}^{\perp})$ have the same rank. Thus if $(G, \boldsymbol{\gamma}, \mathbf{d}^{\perp})$ and hence $(G', \boldsymbol{\gamma}, \mathbf{d}^{\perp})$ have only collapsed realizations, so does $(G', \boldsymbol{\gamma}, \mathbf{d})$ in which case $(G, \boldsymbol{\gamma}, \mathbf{d})$ has no faithful realization.

The Laman direction Now let $(G, \boldsymbol{\gamma})$ be a reflection-Laman graph and let $(G', \boldsymbol{\gamma})$ be a Ross basis of $(G, \boldsymbol{\gamma})$. For any edge $ij \notin G'$, adding it to G' induces a Ross circuit[4] which contains some edge $i'j'$ having the property specified in Proposition 5.3. Note that $G' - ij + i'j'$ is again a Ross basis. We therefore can assume (after edge-swapping in this manner) for all $ij \notin G'$ that ij has the property from Proposition 5.3 in the Ross circuit it induces.

We assign directions $\mathbf{d}'$ to the edges of G' such that:

- The directions on each of the intersections of the Ross circuits with G' are generic in the sense of Proposition 5.3.
- The directions on the edges of G' that remain in the reduced graph $(G^*, \boldsymbol{\gamma})$ are perpendicular to an assignment of directions on G^* that is generic in the sense of Proposition 5.2.
- The directions on the edges of G' are generic in the sense of Proposition 5.1.

This is possible because the set of disallowed directions is the union of a finite number of proper algebraic subsets in the space of direction assignments. Extend to directions $\mathbf{d}$ on G by assigning directions to the remaining edges as specified in Proposition 5.3. By construction, we know that:

Lemma 5.10 *The direction network* $(\tilde{G}, \varphi, \mathbf{d})$ *has faithful realizations.*

Proof The realization space is identified with that of $(\tilde{G}', \varphi, \mathbf{d}')$, and $\mathbf{d}'$ is chosen so that Proposition 5.1 applies. □

Lemma 5.11 *In any realization of* $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$*, the Ross circuits are realized with all their vertices coincident and on the* y*-axis.*

Proof This follows from how we chose $\mathbf{d}$ and Proposition 5.3. □

[4] Recall that here we are using Ross circuit to refer to only one kind of circuit in the Ross matroid. The other type of circuit cannot appear since reflection-Laman graphs do not have (2, 2) blocks with trivial ρ-image.

As a consequence of Lemma 5.11 and the fact that we picked $\mathbf{d}$ so that $\mathbf{d}^{\perp}$ extends to a generic assignment of directions $(\mathbf{d}^{*})^{\perp}$ on the reduced graph $(G^{*}, \boldsymbol{\gamma})$ we have:

Lemma 5.12 *The realization space of* $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ *is identified with that of* $(\tilde{G}^{*}, \varphi, (\mathbf{d}^{*})^{\perp})$ *which, furthermore, contains only collapsed solutions.*

Observe that a direction network for a single self-loop (colored 1) with a generic direction only has solutions where vertices are collapsed and on the y-axis. Consequently, replacing a Ross circuit with a single vertex and a self-loop yields isomorphic realization spaces. Since the reduced graph is reflection-(2, 2) by Proposition 2.6 and the directions assigned to its edges were chosen generically for Proposition 5.2, that $(\tilde{G}, \varphi, \mathbf{d}^{\perp})$ has only collapsed solutions follows. Thus, we have exhibited a special pair, completing the proof.

Remark It can be seen that the realization space of a direction network as supplied in Theorem 4 has at least one degree of freedom for each edge that is not in a Ross basis. Thus, the statement cannot be improved to, e.g., a unique realization up to translation and scale.

6 Infinitesimal Rigidity of Reflection Frameworks

Let $(\tilde{G}, \varphi, \boldsymbol{\ell})$ be a reflection framework, and let $(G, \boldsymbol{\gamma})$ be the quotient graph with n vertices. The algebraic steps in this section are similar to those in Sect. 4. For a reflection framework, the realization space $\mathcal{R}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ defined by (1)–(2) is canonically identified with the solutions to

$$\|\gamma_{ij} \cdot \mathbf{p}_j - \mathbf{p}_i\|^2 = \ell_{ij}^2 \qquad \text{for all edges } ij \in E(G). \tag{16}$$

As in Sect. 5, we assume, without loss of generality, that Γ acts by reflections through the y-axis.

Computing the formal differential of (16), we obtain the system

$$\left\langle \gamma_{ij} \cdot \mathbf{p}_j - \mathbf{p}_i, \mathbf{v}_j - \mathbf{v}_i \right\rangle = 0 \qquad \text{for all edges } ij \in E(G) \tag{17}$$

where the unknowns are the *velocity vectors* $\mathbf{v}_i$. A realization is *infinitesimally rigid* if the system (17) has rank $2n - 1$. As in the case of cone frameworks, generically, infinitesimal rigidity and rigidity coincide, and the non-generic set is defined in the same way.

6.1 Relation to Direction Networks

Here is the core of the direction network method for reflection frameworks: we can understand the rank of (17) in terms of a direction network.

Proposition 6.1 *Let* $\tilde{G}(\mathbf{p})$ *be a realization of a reflection framework. Define the direction* $\mathbf{d}_{ij}$ *to be* $\gamma_{ij} \cdot \mathbf{p}_j - \mathbf{p}_i$. *Then the rank of* (17) *is equal to that of* (13) *for the direction network* $(G, \boldsymbol{\gamma}, \mathbf{d}^{\perp})$.

Proof Exchange the roles of $\mathbf{v}_i$ and $\mathbf{p}_i$ in (17). □

6.2 Proof of Theorem 2

The, more difficult, "Laman direction" of the Main Theorem follows immediately from Theorem 4 and Proposition 6.1: given a reflection-Laman graph, Theorem 4 produces a realization with no coincident endpoints and a certificate that (17) has corank one. The "Maxwell direction" follows from a similar argument as that for Theorem 1.

Remark The statement of Proposition 6.1 is *exactly the same* as the analogous statement for orientation-preserving cases of this theory. What is different is that, for reflection frameworks, the rank of $(G, \boldsymbol{\gamma}, \mathbf{d}^{\perp})$ is *not* the same as that of $(G, \boldsymbol{\gamma}, \mathbf{d})$. By Proposition 5.2, the set of directions arising as the difference vectors from point sets is *always non-generic* on reflection-Laman graphs, so we are forced to introduce the notion of a special pair.

Remark We can extend Theorem 2 to Γ-actions with inverted edges but which are otherwise free. Indeed, a similar argument as in Sect. 4.3 applies here.

Acknowledgments LT was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement No. 247029-SDModels. JM was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement No. 226135. LT and JM had been supported by NSF CDI-I Grant DMR 0835586 to Igor Rivin and M. M. J. Treacy.

References

1. Berardi, M., Heeringa, B., Malestein, J., Theran, L.: Rigid components in fixed-lattice and cone frameworks. In: Proceedings of the 23rd Annual Canadian Conference on Computational Geometry (CCCG) (2011). arXiv:1105.3234
2. Bölcskei, A., Szél-Kopolyás, M.: Construction of D-graphs related to periodic tilings. KoG **6**, 21–27 (2002)
3. Borcea, C., Streinu, I., Tanigawa, S.: Periodic body-and-bar frameworks. SIAM J. Discrete Math. **29**(1), 93–112 (2015). doi:10.1137/120900265
4. Borcea, C.S., Streinu, I.: Periodic frameworks and flexibility. Proc. R. Soc. Lond. Ser. A **466**(2121), 2633–2649 (2010). doi:10.1098/rspa.2009.0676
5. Borcea, C.S., Streinu, I.: Minimally rigid periodic graphs. Bull. Lond. Math. Soc. **43**(6), 1093–1103 (2011). doi:10.1112/blms/bdr044
6. Brylawski, T.: Constructions. Theory of Matroids. Encyclopedia of Mathematics and Its Applications, vol. 26, pp. 127–223. Cambridge University Press, Cambridge (1986). doi:10.1017/CBO9780511629563.010
7. Develin, M., Martin, J.L., Reiner, V.: Rigidity theory for matroids. Comment. Math. Helv. **82**(1), 197–233 (2007). doi:10.4171/CMH/89
8. Edmonds, J., Rota, G.-C.: Submodular set functions (abstract). In: Waterloo Combinatorics Conference, University of Waterloo, Ontario (1966)
9. Fowler, P.W., Guest, S.D.: A symmetry extension of Maxwell's rule for rigidity of frames. Int. J. Solids Struct. **37**(12), 1793–1804 (1999)
10. Jordán, T., Kaszanitzky, V., Tanigawa, S.: Gain-sparsity and symmetric rigidity in the plane. Technical Report (2012). http://www.cs.elte.hu/egres/tr/egres-12-17
11. Király, F., Theran, L., Tomioka, R.: The algebraic and combinatorial approach to matrix completion. To appear in J. Mach. Learn. Res. (2015). arXiv:1211.4116. With an appendix by T. Uno
12. Laman, G.: On graphs and rigidity of plane skeletal structures. J. Eng. Math. **4**, 331–340 (1970)

13. Lee, A., Streinu, I.: Pebble game algorithms and sparse graphs. Discrete Math. **308**(8), 1425–1437 (2008). doi:10.1016/j.disc.2007.07.104
14. Malestein, J., Theran, L.: Generic rigidity of frameworks with orientation-preserving crystallographic symmetry. Preprint, arXiv:1108.2518 (2011)
15. Malestein, J., Theran, L.: Generic rigidity of reflection frameworks. Preprint, arXiv:1203.2276 (2012)
16. Malestein, J., Theran, L.: Generic combinatorial rigidity of periodic frameworks. Adv. Math. **233**, 291–331 (2013). doi:10.1016/j.aim.2012.10.007
17. Malestein, J., Theran, L.: Generic rigidity with forced symmetry and sparse colored graphs. In: Connelly, R., Weiss, A.I., Whiteley, W. (eds.) Rigidity and Symmetry. Fields Institute Communications, vol. 70, pp. 227–252. Springer, New York (2014). doi:10.1007/978-1-4939-0781-6_12
18. Malestein, J., Theran, L.: Frameworks with forced symmetry II: orientation-preserving crystallographic groups. Geom. Dedicata **170**, 219–262 (2014). doi:10.1007/s10711-013-9878-6
19. Maxwell, J.C.: On the calculation of the equilibrium and stiffness of frames. Philos. Mag. **27**, 294 (1864)
20. Milnor, J.: Singular points of complex hypersurfaces. Annals of Mathematics Studies, vol. 61. Princeton University Press, Princeton (1968)
21. Nixon, A., Ross, E.: Periodic rigidity on a variable torus using inductive constructions (2012). arXiv:1204.1349
22. Rivin, I.: Geometric simulations: a lesson from virtual zeolites. Nat. Mater. **5**, 931–932 (2006). doi:10.1038/nmat1792
23. Ross, E.: Inductive constructions for frameworks on a two-dimensional fixed torus. Preprint, arXiv:1203.6561 (2012)
24. Ross, E., Schulze, B., Whiteley, W.: Finite motions from periodic frameworks with added symmetry. Int. J. Solids Struct. **48**(11–12), 1711–1729 (2011). doi:10.1016/j.ijsolstr.2011.02.018
25. Sartbaeva, A., Wells, S., Treacy, M., Thorpe, M.: The flexibility window in zeolites. Nat. Mater. **5**(12), 962–965 (2006)
26. Schulze, B.: Symmetric Laman theorems for the groups $\mathcal{C}_2$ and $\mathcal{C}_s$. Electron. J. Comb. *17*(1): Research Paper 154, 61 (2010). http://www.combinatorics.org/Volume_17/Abstracts/v17i1r154.html
27. Schulze, B.: Symmetric versions of Laman's theorem. Discrete Comput. Geom. **44**(4), 946–972 (2010). doi:10.1007/s00454-009-9231-x
28. Schulze, B., Tanigawa, S.I.: Linking rigid bodies symmetrically. Eur. J. Comb. **42**(0), 145–166 (2014). doi:10.1016/j.ejc.2014.06.002
29. Schuze, B., Tanigawa, S.I.: Infinitesimal rigidity of symmetric frameworks. Preprint, arXiv: 1308.6380 (2013)
30. Streinu, I., Theran, L.: Sparsity-certifying graph decompositions. Graphs Comb. **25**(2), 219–238 (2009). doi:10.1007/s00373-008-0834-4
31. Streinu, I., Theran, L.: Slider-pinning rigidity: a Maxwell-Laman-type theorem. Discrete Comput. Geom. **44**(4), 812–837 (2010). doi:10.1007/s00454-010-9283-y
32. Tanigawa, S-I.: Matroids of gain graphs in applied discrete geometry. Preprint, arXiv:1207.3601 (2012)
33. Treacy, M.M.J., Foster, M.D., Rivin, I.: Towards a catalog of designer zeolites. Turning Points in Solid State. Materials and Surface Science, pp. 208–220. RSC Publishing, Cambridge (2008)
34. Whiteley, W.: The union of matroids and the rigidity of frameworks. SIAM J. Discrete Math. **1**(2), 237–255 (1988). doi:10.1137/0401025
35. Whiteley, W.: Some matroids from discrete applied geometry. In: Bonin, J., Oxley, J.G., Servatius, B. (eds.) Matroid Theory. Contemporary Mathematics, vol. 197, pp. 171–311. American Mathematical Society, Providence, RI (1996)
36. Zaslavsky, T.: Voltage-graphic matroids. Matroid Theory and Its Applications, pp. 417–424. Liguori, Naples (1982)