Rigidity of symmetric frameworks in normed spaces

D. Kitson[1,*]

Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, U.K.

Department of Mathematics and Computer Studies, Mary Immaculate College, Thurles, Co. Tipperary, Ireland.

A. Nixon[*], B. Schulze[*]

Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, U.K.

Abstract

We develop a combinatorial rigidity theory for symmetric bar-joint frameworks in a general finite dimensional normed space. In the case of rotational symmetry, matroidal Maxwell-type sparsity counts are identified for a large class of d-dimensional normed spaces (including all ℓ^p spaces with $p \neq 2$). Complete combinatorial characterisations are obtained for half-turn rotation in the ℓ^1 and ℓ^∞-plane. As a key tool, a new Henneberg-type inductive construction is developed for the matroidal class of $(2, 2, 0)$-gain-tight gain graphs.

Keywords: bar-joint framework, infinitesimal rigidity, gain graphs, matroids, normed spaces.
2010 MSC: 52C25, 05C50, 05B35

Contents

1 Introduction **2**

2 Symmetric frameworks and gain sparsity **3**

3 An inductive construction of $(2, 2, 0)$-gain tight $\mathbb{Z}_2$-gain graphs **15**

4 Application to C_2-symmetric frameworks in the ℓ^1 and ℓ^∞-plane **37**

5 Concluding remarks **48**

[*]Corresponding author; phone: +353 504 20392.
Email addresses: `derek.kitson@mic.ul.ie` (D. Kitson), `a.nixon@lancaster.ac.uk` (A. Nixon), `b.schulze@lancaster.ac.uk` (B. Schulze)

[1]Supported by the Engineering and Physical Sciences Research Council [grant numbers EP/P01108X/1 and EP/S00940X/1].

1. Introduction

The determination of rigidity and flexibility of bar-joint frameworks consisting of rigid bars connected at their ends by idealised joints is a highly active research area in discrete geometry with a long and rich history dating back to considerations of linkages, trusses and polyhedral structures by Maxwell, Cauchy and Euler, among others. Since bar-joint frameworks are suitable models for a variety of both man-made and natural structures (buildings, linkages, molecules, crystals, etc.), rigidity theory has a broad range of modern practical applications in fields such as engineering, robotics, CAD and materials science. (See [18, 19, 20] e.g.). This transfer of knowledge between fundamental and applied researchers is one of the motivations for exploring constraint systems in new geometric contexts, such as the normed spaces considered in this article (see also [1, 2] for related problems). Another strong motivation comes from the potential for developing combinatorial Laman-type characterisations ([8]) of rigid graphs in any dimension, due to the amenability of the matroidal sparsity counts arising in some of these contexts.

In this article, we consider first-order rigidity for bar-joint frameworks with a finite abelian symmetry group, developing both a general linear theory as well as sharp combinatorial results in the case of half-turn rotational symmetry in the ℓ^1 and ℓ^∞-plane. This complements and extends work of Schulze [16], Jordán, Kaszanitzky and Tanigawa [4], Malestein and Theran [11], and Schulze and Tanigawa [17] on symmetric frameworks in Euclidean space, and work of Kitson and Power [5] and Kitson and Schulze [6, 7] on infinitesimal rigidity in normed spaces.

In Sect. 2, we introduce the natural notion of a *framework complex* and develop several key tools for analysing frameworks with a finite abelian symmetry group acting freely on the vertex set. These include a decomposition theorem for the framework complex (which incorporates a block decomposition for the rigidity operator) and counting criteria for the accompanying group-labelled quotient graph (called a *gain graph*). For a large class of d-dimensional normed spaces, with $d \geq 2$, this leads to the identification of (d, d, m)-gain-tight gain graphs, with $m \in \{0, 1, 2, d-2\}$, as the underlying structure graphs for phase-symmetrically isostatic frameworks with rotational symmetry (see Corollary 2.20). In contrast to Euclidean contexts, these classes of graphs are matroidal for all dimensions d, and are computationally accessible through associated pebble-game algorithms [9].

In Sect. 3, a new inductive construction is obtained for the class of $(2, 2, 0)$-gain-tight gain graphs (Theorem 3.20). Previous recursive characterisations of $(2, 2, m)$-gain-tight graphs, with $m \in \{1, 2\}$, can be found in [15]. The construction presented here is necessarily more involved due to a step change in the possible minimum degree when $m = 0$. Recursive constructions of classes of graphs are of fundamental importance in rigidity theory, occurring for example in Laman's landmark characterisation of rigidity in the Euclidean plane [8]. Of particular relevance are previous characterisations of classes of gain graphs [4, 15] and characterisations where graph simplicity is required to be preserved [12, 13].

In Sect. 4, geometric and combinatorial characterisations are obtained for the rigidity of two-dimensional frameworks with half-turn rotational symmetry in the ℓ^1 and ℓ^∞-plane. The geometric results (Theorem 4.2) use an edge-colouring technique which expresses the gain graph of a phase-symmetrically isostatic framework as an edge-disjoint union of either two unbalanced spanning map graphs (defined in Sect. 4.1), or two spanning trees. Combinatorial characterisations are then

obtained for graphs which admit a placement as a phase-symmetrically isostatic framework with half-turn rotational symmetry (Theorems 4.3 and 4.9) by combining these geometric results with the construction scheme from Sect. 3. The analogous problem for frameworks with reflectional symmetry requires different methods and was settled in [7].

2. Symmetric frameworks and gain sparsity

The aim of this section is to derive necessary gain-graph counting conditions for symmetrically isostatic bar-joint frameworks in normed spaces. Throughout this article, X denotes a finite dimensional real vector space with a norm $\|\cdot\|$ and dimension $d \geq 2$. The group of linear isometries of X is denoted $\mathrm{Isom}(X, \|\cdot\|)$, or simply $\mathrm{Isom}(X)$. The complexification $\mathbb{C} \otimes_{\mathbb{R}} X$ is denoted $X_{\mathbb{C}}$ and, for convenience, elementary tensor products of the form $\lambda \otimes x$ are denoted by λx. Also, Γ will denote a finite abelian group with identity element 1 and $\hat{\Gamma}$ will denote the dual group of characters $\chi : \Gamma \to \{z \in \mathbb{C} : |z| = 1\}$.

2.1. Bar-joint frameworks

Let $G = (V, E)$ be a finite simple undirected graph and let $p = (p_v)_{v \in V} \in X^V$. If the components of p are distinct vectors in X then the pair (G, p) is called a *bar-joint framework* in X. If H is a subgraph of G and $p_H = (p_v)_{v \in V(H)}$ then the pair (H, p_H) is called a *subframework* of (G, p). Define,

$$f_G : X^V \to \mathbb{R}^E, \qquad (x_v)_{v \in V} \mapsto (\|x_v - x_w\|)_{vw \in E}.$$

If f_G is differentiable at p then the bar-joint framework (G, p) is said to be *well-positioned* in X.

Lemma 2.1. *[6, Proposition 6] If (G, p) is well-positioned in X then the differential of f_G at p satisfies,*

$$df_G(p) : X^V \to \mathbb{R}^E, \quad (u_v)_{v \in V} \mapsto (\varphi_{v,w}(u_v - u_w))_{vw \in E},$$

where, for each edge $vw \in E$, $\varphi_{v,w} : X \to \mathbb{R}$ is the linear functional,

$$\varphi_{v,w}(x) = \lim_{t \to 0} \frac{1}{t} \left(\|p_v - p_w + tx\| - \|p_v - p_w\| \right), \quad \forall\, x \in X.$$

A *rigid motion* of X is a family of continuous paths $\{\alpha_x : [-1, 1] \to X\}_{x \in X}$ such that $\alpha_x(0) = x$ and $\|\alpha_x(t) - \alpha_y(t)\| = \|x - y\|$ for all pairs $x, y \in X$ and all $t \in [-1, 1]$. An *infinitesimal rigid motion* of X is a vector field $\eta : X \to X$ with the property that $\eta(x) = \alpha'_x(0)$ for all $x \in X$, for some rigid motion $\{\alpha_x\}_{x \in X}$. The collection of all infinitesimal rigid motions of X is a vector subspace of X^X, denoted $\mathcal{T}(X)$.

Let (G, p) be a bar-joint framework in X and define $\rho_{(G,p)} : \mathcal{T}(X) \to X^V$, $\eta \mapsto (\eta(p_v))_{v \in V}$. Note that if (G, p) is well-positioned, then $df_G(p) \circ \rho_{(G,p)} = 0$ (see [5, Lemma 2.1]). The *framework complex* for (G, p), denoted $\mathrm{comp}(G, p)$, is the chain complex,

$$0 \longrightarrow \mathcal{T}(X) \xrightarrow{\rho_{(G,p)}} X^V \xrightarrow{df_G(p)} \mathbb{R}^E \longrightarrow 0. \tag{1}$$

The kernel of $df_G(p)$, denoted $\mathcal{F}(G, p)$, is referred to as the space of *infinitesimal flexes* of (G, p), while the image of $\rho_{(G,p)}$, denoted $\mathcal{T}(G, p)$, is referred to as the space of *trivial* infinitesimal flexes of (G, p).

Definition 2.2. A well-positioned bar-joint framework (G, p) in a normed space X is,

(a) *full* if the framework complex $\mathrm{comp}(G, p)$ is exact at $\mathcal{T}(X)$.

(b) *infinitesimally rigid* if the framework complex $\mathrm{comp}(G, p)$ is exact at X^V.

(c) *independent* if the framework complex $\mathrm{comp}(G, p)$ is exact at $\mathbb{R}^E$.

A well-positioned bar-joint framework is *isostatic* if it is both infinitesimally rigid and independent. Note that $\mathrm{comp}(G, p)$ is a short exact sequence if and only if (G, p) is both full and isostatic.

2.2. Symmetric graphs

A Γ-*symmetric graph* is a pair (G, θ) where G is a finite simple undirected graph with automorphism group $\mathrm{Aut}(G)$ and $\theta : \Gamma \to \mathrm{Aut}(G)$ is a group homomorphism. It is assumed throughout this article that θ acts freely on the vertex set of G. Thus $v \neq \theta(\gamma)v$ for all $v \in V$ and for all $\gamma \in \Gamma$ with $\gamma \neq 1$. For convenience, we suppress θ and denote $\theta(\gamma)$ by γ for each group element $\gamma \in \Gamma$. Also, for each edge $e = vw \in E$ we denote by γe, or $\gamma(vw)$, the edge in E which joins the vertices γv and γw.

Proposition 2.3. *Let (G, θ) be a Γ-symmetric graph and let $\tau : \Gamma \to \mathrm{Isom}(X)$ be a group representation.*

(i) $(X_{\mathbb{C}})^V = \bigoplus_{\chi \in \hat{\Gamma}} X_\chi$ *where,*

$$X_\chi = \{x = (x_v)_{v \in V} \in (X_{\mathbb{C}})^V : x_{\gamma v} = \chi(\gamma)\tau(\gamma)x_v, \ \forall\ \gamma \in \Gamma, \ \forall\ v \in V\}.$$

(ii) $\mathbb{C}^E = \bigoplus_{\chi \in \hat{\Gamma}} Y_\chi$ *where,*

$$Y_\chi = \{y = (y_e)_{e \in E} \in \mathbb{C}^E : y_{\gamma e} = \chi(\gamma)y_e, \ \forall\ \gamma \in \Gamma, \ \forall\ e \in E\}.$$

Proof. Each $x = (x_v)_{v \in V} \in (X_{\mathbb{C}})^V$ may be expressed as a sum $x = \sum_{\chi \in \hat{\Gamma}} x_\chi$ where $x_\chi = (x_{\chi,v})_{v \in V} \in (X_{\mathbb{C}})^V$ has components,

$$x_{\chi,v} = \frac{1}{|\Gamma|}\left(\sum_{\gamma \in \Gamma} \overline{\chi(\gamma)}\tau(\gamma^{-1})x_{\gamma v}\right).$$

Similarly, each $y = (y_e)_{e \in E} \in \mathbb{C}^E$ may be expressed as a sum $\sum_{\chi \in \hat{\Gamma}} y_\chi$ where $y_\chi = (y_{\chi,e})_{e \in E} \in \mathbb{C}^E$ has components,

$$y_{\chi,e} = \frac{1}{|\Gamma|}\left(\sum_{\gamma \in \Gamma} \overline{\chi(\gamma)}y_{\gamma e}\right).$$

We use here the standard identity,

$$\sum_{\chi \in \hat{\Gamma}} \chi(\gamma) = \begin{cases} |\Gamma| & \text{if } \gamma = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $x_\chi \in X_\chi$ and $y_\chi \in Y_\chi$ for each $\chi \in \hat{\Gamma}$.

Let $S \subseteq \hat{\Gamma}$ be a maximal subset which gives rise to a direct sum $\oplus_{\chi\in S} X_\chi$ and suppose there exists $\tilde{\chi} \in \hat{\Gamma}\backslash S$. If $x \in X_{\tilde{\chi}} \cap (\oplus_{\chi\in S} X_\chi)$ then $x_v = \overline{\tilde{\chi}(\gamma)}\tau(\gamma^{-1})x_{\gamma v}$ for all $v \in V$ and all $\gamma \in \Gamma$. Moreover, $x = \sum_{\chi\in S} z_\chi$ for some unique $z_\chi \in X_\chi$. It follows that, for all $v \in V$ and all $\gamma \in \Gamma$,

$$\begin{aligned} x_v &= \overline{\tilde{\chi}(\gamma)}\tau(\gamma^{-1})\left(\sum_{\chi\in S} z_{\chi,\gamma v}\right) \\ &= \overline{\tilde{\chi}(\gamma)}\tau(\gamma^{-1})\left(\sum_{\chi\in S} \chi(\gamma)\tau(\gamma)z_{\chi,v}\right) \\ &= \overline{\tilde{\chi}(\gamma)}\left(\sum_{\chi\in S} \chi(\gamma)z_{\chi,v}\right). \end{aligned}$$

Thus $x = \overline{\tilde{\chi}(\gamma)}(\sum_{\chi\in S} \chi(\gamma)z_\chi)$ for all $\gamma \in \Gamma$. Since the sum $x = \sum_{\chi\in S} z_\chi$ is direct, if $x \neq 0$ then $\tilde{\chi} = \chi$ for some $\chi \in S$. This is a contradiction and so $X_{\tilde{\chi}} \cap (\oplus_{\chi\in S} X_\chi) = \{0\}$. However, this contradicts the maximality of S and so it follows that $S = \hat{\Gamma}$. This establishes the direct sum $\oplus_{\chi\in\hat{\Gamma}} X_\chi$ and a similar argument can be applied for $\oplus_{\chi\in\hat{\Gamma}} Y_\chi$. □

2.3. Symmetric frameworks

A Γ-*symmetric bar-joint framework* is a tuple $\mathcal{G} = (G, p, \theta, \tau)$ where (G, p) is a bar-joint framework, (G, θ) is a Γ-symmetric graph and $\tau : \Gamma \to \mathrm{Isom}(X)$ is a group representation which satisfies $\tau(\gamma)(p_v) = p_{\gamma v}$ for all $\gamma \in \Gamma$ and all $v \in V$.

Lemma 2.4. *Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X and let $vw \in E$. Then $\varphi_{\gamma v,\gamma w} = \varphi_{v,w} \circ \tau(\gamma^{-1})$ for all $\gamma \in \Gamma$.*

Proof. Let $p_0 = p_v - p_w$. Then $\tau(\gamma)(p_0) = p_{\gamma v} - p_{\gamma w}$ and so for each $x \in X$,

$$\begin{aligned} \varphi_{\gamma v,\gamma w}(x) &= \lim_{t\to 0} \frac{1}{t}\left(\|\tau(\gamma)(p_0 + t\,\tau(\gamma^{-1})x)\| - \|\tau(\gamma)p_0\|\right) \\ &= \lim_{t\to 0} \frac{1}{t}\left(\|p_0 + t\,\tau(\gamma^{-1})x\| - \|p_0\|\right) = \varphi_{v,w}(\tau(\gamma^{-1})x). \end{aligned}$$

□

In the following, the same symbol will be used to denote a real affine transformation $T : Y \to Z$ between two real linear spaces Y and Z and its complex extension $T : Y_{\mathbb{C}} \to Z_{\mathbb{C}}$. In particular, we consider the complex linear functionals $\varphi_{v,w} : X_{\mathbb{C}} \to \mathbb{C}$, the complex linear transformations $\tau(\gamma) : X_{\mathbb{C}} \to X_{\mathbb{C}}$ and the complex differential $df_G(p) : (X_{\mathbb{C}})^V \to \mathbb{C}^E$ associated to a bar-joint framework (G, p) in X.

Proposition 2.5. *Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X. With respect to the direct sum decompositions obtained in Proposition 2.3,*

$$(X_{\mathbb{C}})^V = \bigoplus_{\chi\in\hat{\Gamma}} X_\chi \quad \textit{and} \quad \mathbb{C}^E = \bigoplus_{\chi\in\hat{\Gamma}} Y_\chi,$$

the (complex) differential $df_G(p)$ may be expressed as a direct sum of linear transformations,

$$df_G(p) = \bigoplus_{\chi\in\hat{\Gamma}} R_\chi(\mathcal{G}),$$

where $R_\chi(\mathcal{G}) : X_\chi \to Y_\chi$ for each character $\chi \in \hat{\Gamma}$.

Proof. Let $vw \in E$. If $(x_v)_{v\in V} \in X_\chi$ then, using Lemma 2.4,

$$\varphi_{v,w}(x_v - x_w) = \varphi_{\gamma v,\gamma w}(\tau(\gamma)(x_v - x_w)) = \varphi_{\gamma v,\gamma w}(\overline{\chi(\gamma)}(x_{\gamma v} - x_{\gamma w})) = \overline{\chi(\gamma)}\varphi_{\gamma v,\gamma w}(x_{\gamma v} - x_{\gamma w}),$$

for each $\chi \in \hat{\Gamma}$ and each $\gamma \in \Gamma$. Thus, by Lemma 2.1, $df_G(p)(X_\chi) \subseteq Y_\chi$ and the result follows. □

2.4. Infinitesimal rigid motions

Denote by $\mathcal{T}(X;\mathbb{C})$ the complex vector space spanned by vector fields $\eta_\mathbb{C} : X \to X_\mathbb{C}$, $x \mapsto 1 \otimes \eta(x)$ where $\eta \in \mathcal{T}(X)$. For convenience, $\eta_\mathbb{C}$ will simply be denoted η.

Proposition 2.6. *Let $\tau : \Gamma \to \mathrm{Isom}(X)$ be a group representation. Then, $\mathcal{T}(X;\mathbb{C}) = \bigoplus_{\chi\in\hat{\Gamma}} \mathcal{T}_\chi(X)$ where,*

$$\mathcal{T}_\chi(X) = \{\eta \in \mathcal{T}(X;\mathbb{C}) : \eta(\tau(\gamma)x) = \chi(\gamma)\tau(\gamma)\eta(x),\ \forall\ \gamma \in \Gamma,\ \forall\ x \in X\}.$$

Proof. Applying an argument similar to Lemma 2.3, each $\eta \in \mathcal{T}(X;\mathbb{C})$ may be expressed as a sum $\eta = \sum_{\chi\in\hat{\Gamma}} \eta_\chi$ where $\eta_\chi : X \to X_\mathbb{C}$ is the vector field,

$$\eta_\chi(x) = \frac{1}{|\Gamma|}\left(\sum_{\gamma\in\Gamma} \overline{\chi(\gamma)}\tau(\gamma^{-1})\eta(\tau(\gamma)x)\right).$$

Note that $\eta_\chi \in \mathcal{T}_\chi(X)$ for each $\chi \in \hat{\Gamma}$.

Let $S \subseteq \hat{\Gamma}$ be a maximal subset which gives rise to a direct sum $\oplus_{\chi\in S}\mathcal{T}_\chi(X)$ and suppose there exists $\tilde{\chi} \in \hat{\Gamma}\backslash S$. If $\eta \in \mathcal{T}_{\tilde{\chi}}(X) \cap (\oplus_{\chi\in S}\mathcal{T}_\chi(X))$ then $\eta(x) = \overline{\tilde{\chi}(\gamma)}\tau(\gamma^{-1})\eta(\tau(\gamma)x)$ for all $x \in X$ and all $\gamma \in \Gamma$. Moreover, $\eta = \sum_{\chi\in S} \delta_\chi$ for some unique $\delta_\chi \in \mathcal{T}_\chi(X)$. Thus, for all $x \in X$ and all $\gamma \in \Gamma$,

$$\begin{aligned}
\eta(x) &= \overline{\tilde{\chi}(\gamma)}\tau(\gamma^{-1})\left(\sum_{\chi\in S} \delta_\chi(\tau(\gamma)x)\right) \\
&= \overline{\tilde{\chi}(\gamma)}\tau(\gamma^{-1})\left(\sum_{\chi\in S} \chi(\gamma)\tau(\gamma)\delta_\chi(x)\right) \\
&= \overline{\tilde{\chi}(\gamma)}\left(\sum_{\chi\in S} \chi(\gamma)\delta_\chi(x)\right).
\end{aligned}$$

Thus $\eta = \overline{\tilde{\chi}(\gamma)}\left(\sum_{\chi\in S} \chi(\gamma)\delta_\chi\right)$ for all $\gamma \in \Gamma$. Since the sum $\eta = \sum_{\chi\in S} \delta_\chi$ is direct, if $\eta \neq 0$ then it follows that $\tilde{\chi} = \chi$ for some $\chi \in S$. This is a contradiction and so $\mathcal{T}_{\tilde{\chi}}(X) \cap (\oplus_{\chi\in S}\mathcal{T}_\chi(X)) = \{0\}$. However, this contradicts the maximality of S and so $S = \hat{\Gamma}$.

□

We now consider the complex restriction map $\rho_{(G,p)} : \mathcal{T}(X;\mathbb{C}) \to (X_{\mathbb{C}})^V$, $\eta \mapsto (\eta(p_v))_{v\in V}$ for a given bar-joint framework (G,p) in X.

Proposition 2.7. *Let $\mathcal{G} = (G,p,\theta,\tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X. With respect to the direct sum decompositions obtained in Propositions 2.6 and 2.3,*

$$\mathcal{T}(X;\mathbb{C}) = \bigoplus_{\chi\in\hat{\Gamma}} \mathcal{T}_\chi(X), \quad \textit{and,} \quad (X_{\mathbb{C}})^V = \bigoplus_{\chi\in\hat{\Gamma}} X_\chi,$$

the (complex) restriction map $\rho_{(G,p)}$ may be expressed as a direct sum of linear transformations,

$$\rho_{(G,p)} = \bigoplus_{\chi\in\hat{\Gamma}} \rho_\chi(\mathcal{G}),$$

where $\rho_\chi(\mathcal{G}) : \mathcal{T}_\chi(X) \to X_\chi$ for each character $\chi \in \hat{\Gamma}$.

Proof. Let $\chi \in \hat{\Gamma}$. If $\eta \in \mathcal{T}_\chi(X)$ then, for each $\gamma \in \Gamma$ and all $v \in V$,

$$\eta(p_{\gamma v}) = \eta(\tau(\gamma)p_v) = \chi(\gamma)\tau(\gamma)\eta(p_v).$$

Thus, $\rho_{(G,p)}(\mathcal{T}_\chi(X)) \subseteq X_\chi$ and the result follows. □

2.5. Decomposing the framework complex

Denote by $\mathrm{comp}_{\mathbb{C}}(G,p)$ the *complexified* framework complex for a bar-joint framework (G,p),

$$0 \longrightarrow \mathcal{T}(X;\mathbb{C}) \xrightarrow{\rho_{(G,p)}} (X_{\mathbb{C}})^V \xrightarrow{df_G(p)} \mathbb{C}^E \longrightarrow 0. \tag{2}$$

If $\mathcal{G} = (G,p,\theta,\tau)$ is a well-positioned and Γ-symmetric bar-joint framework in X then, recalling the decompositions $df_G(p) = \bigoplus_{\chi\in\hat{\Gamma}} R_\chi(\mathcal{G})$ and $\rho_{(G,p)} = \bigoplus_{\chi\in\hat{\Gamma}} \rho_\chi(\mathcal{G})$ from Propositions 2.5 and 2.7, we have $R_\chi(\mathcal{G}) \circ \rho_\chi(\mathcal{G}) = 0$ for all $\chi \in \hat{\Gamma}$. The *χ-symmetric framework complex* for $\mathcal{G}$, denoted $\mathrm{comp}_\chi(\mathcal{G})$, is the chain complex,

$$0 \longrightarrow \mathcal{T}_\chi(X) \xrightarrow{\rho_\chi(\mathcal{G})} X_\chi \xrightarrow{R_\chi(\mathcal{G})} Y_\chi \longrightarrow 0. \tag{3}$$

The kernel of $R_\chi(\mathcal{G})$, denoted $\mathcal{F}_\chi(\mathcal{G})$, is referred to as the space of *χ-symmetric infinitesimal flexes* of $\mathcal{G}$. The image of ρ_χ, denoted $\mathcal{T}_\chi(\mathcal{G})$, is referred to as the space of *trivial* χ-symmetric infinitesimal flexes of $\mathcal{G}$.

Theorem 2.8. *Let $\mathcal{G} = (G,p,\theta,\tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X. Then,*

$$\mathrm{comp}_{\mathbb{C}}(G,p) = \bigoplus_{\chi\in\hat{\Gamma}} \mathrm{comp}_\chi(\mathcal{G}).$$

Proof. The result follows from Propositions 2.5, 2.6 and 2.7. □

Definition 2.9. A well-positioned and Γ-symmetric bar-joint framework $\mathcal{G} = (G,p,\theta,\tau)$ in a normed space X is said to be,

(a) *χ-symmetrically full* if $\mathrm{comp}_\chi(\mathcal{G})$ is exact at $\mathcal{T}_\chi(X)$.

(b) *χ-symmetrically infinitesimally rigid* if $\mathrm{comp}_\chi(\mathcal{G})$ is exact at X_χ.

(c) *χ-symmetrically independent* if $\mathrm{comp}_\chi(\mathcal{G})$ is exact at Y_χ.

A Γ-symmetric bar-joint framework is *χ-symmetrically isostatic* if it is both χ-symmetrically infinitesimally rigid and χ-symmetrically independent.

Corollary 2.10. *Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X. The following statements are equivalent.*

(i) (G, p) is full (respectively, infinitesimally rigid, independent or isostatic).

(ii) $\mathcal{G}$ is χ-symmetrically full (respectively, χ-symmetrically infinitesimally rigid, χ-symmetrically independent or χ-symmetrically isostatic) for each $\chi \in \hat{\Gamma}$.

Let $\mathcal{G} = (G, p, \theta, \tau)$ be a Γ-symmetric bar-joint framework in X. A Γ-*symmetric subframework* of $\mathcal{G}$ is a Γ-symmetric framework $\mathcal{H} = (H, p_H, \theta_H, \tau_H)$ where (H, p_H) is a subframework of (G, p), $\theta_H : \Gamma \to \mathrm{Aut}(H)$ is the group homomorphism induced by θ and $\tau_H = \tau$.

Lemma 2.11. *Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X and let $\chi \in \hat{\Gamma}$. If $\mathcal{G}$ is χ-symmetrically independent then every Γ-symmetric subframework of $\mathcal{G}$ is χ-symmetrically independent.*

Proof. Let $\mathcal{H} = (H, p_H, \theta_H, \tau_H)$ be a Γ-symmetric subframework of $\mathcal{G}$ and consider the direct sum decompositions,

$$X_{\mathbb{C}}^{V(G)} = X_{\mathbb{C}}^{V(H)} \oplus X_{\mathbb{C}}^{V(G)\setminus V(H)}, \quad \text{and,} \quad \mathbb{C}^{E(G)} = \mathbb{C}^{E(H)} \oplus \mathbb{C}^{E(G)\setminus E(H)}$$

Note that $X_\chi = X_\chi^{\mathcal{H}} \oplus X_\chi^{\mathcal{G}\setminus\mathcal{H}}$ where $X_\chi^{\mathcal{H}} = X_{\mathbb{C}}^{V(H)} \cap X_\chi$ and $X_\chi^{\mathcal{G}\setminus\mathcal{H}} = X_{\mathbb{C}}^{V(G)\setminus V(H)} \cap X_\chi$. Similarly, $Y_\chi = Y_\chi^{\mathcal{H}} \oplus Y_\chi^{\mathcal{G}\setminus\mathcal{H}}$ where $Y_\chi^{\mathcal{H}} = \mathbb{C}^{E(H)} \cap Y_\chi$ and $Y_\chi^{\mathcal{G}\setminus\mathcal{H}} = \mathbb{C}^{E(G)\setminus E(H)} \cap Y_\chi$. With respect to these decompositions, $R_\chi(\mathcal{G})$ admits a block decomposition of the form,

$$R_\chi(\mathcal{G}) = \begin{pmatrix} R_\chi(\mathcal{H}) & 0 \\ C & D \end{pmatrix}.$$

Thus, if $R_\chi(\mathcal{G})$ is surjective then so too is $R_\chi(\mathcal{H})$. □

2.6. Quotient graphs

Let (G, θ) be a Γ-symmetric graph and suppose θ acts freely on the vertices and edges of G. The orbit of a vertex $v \in V$ (respectively an edge $e \in E$) is denoted by $[v]$ (respectively $[e]$). Thus $[v] = \{\gamma v : \gamma \in \Gamma\}$ and $[e] = \{\gamma e : \gamma \in \Gamma\}$. The collection of all vertex orbits (respectively, edge orbits) is denoted V_0 (respectively, E_0). The *quotient graph* $G_0 = G/\Gamma$ is a multigraph with vertex set V_0, edge set E_0 and incidence relation satisfying $[e] = [v][w]$ if some (equivalently, every) edge in $[e]$ is incident with a vertex in $[v]$ and a vertex in $[w]$.

Proposition 2.12. *Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X. Let $\chi \in \hat{\Gamma}$ and suppose $\mathcal{G}$ is χ-symmetrically full.*

(i) If $\mathcal{G}$ is χ-symmetrically infinitesimally rigid then,

$$|E_0| \geq (\dim_{\mathbb{R}} X)|V_0| - \dim_{\mathbb{C}} \mathcal{T}_\chi(X).$$

(ii) If $\mathcal{G}$ is χ-symmetrically independent then,

$$|E_0| \leq (\dim_{\mathbb{R}} X)|V_0| - \dim_{\mathbb{C}} \mathcal{T}_\chi(X).$$

(iii) If $\mathcal{G}$ is χ-symmetrically isostatic then,

$$|E_0| = (\dim_{\mathbb{R}} X)|V_0| - \dim_{\mathbb{C}} \mathcal{T}_\chi(X).$$

Proof. Applying Proposition 2.5,

$$|E_0| = \dim_{\mathbb{C}} Y_\chi \geq \operatorname{rank} R_\chi = \dim_{\mathbb{C}} X_\chi - \dim_{\mathbb{C}} \ker R_\chi = (\dim_{\mathbb{C}} X_{\mathbb{C}})|V_0| - \dim_{\mathbb{C}} \mathcal{F}_\chi(\mathcal{G}).$$

If (i) holds then $\mathcal{F}_\chi(\mathcal{G}) = \mathcal{T}_\chi(\mathcal{G})$ and $\dim_{\mathbb{C}} \mathcal{T}_\chi(\mathcal{G}) = \dim_{\mathbb{C}} \mathcal{T}_\chi(X)$. If (ii) holds then $\dim_{\mathbb{C}} Y_\chi = \operatorname{rank} R_\chi$ and $\dim_{\mathbb{C}} \mathcal{F}_\chi(\mathcal{G}) \geq \dim_{\mathbb{C}} \mathcal{T}_\chi(\mathcal{G}) = \dim_{\mathbb{C}} \mathcal{T}_\chi(X)$. If (iii) holds then the result follows from (i) and (ii). □

2.7. Norms with a minimal space of infinitesimal rigid motions

The space $\mathcal{T}(X)$ of infinitesimal rigid motions of a normed space X is *minimal* if given any $\eta \in \mathcal{T}(X)$ there exists $x_0 \in X$ such that $\eta(x) = x_0$ for all $x \in X$. This class includes all ℓ^p-spaces, with $p \neq 2$, and all normed spaces with a polyhedral unit ball (see [5, Lemma 2.5]). If $\dim_{\mathbb{R}} X = 2$, then this class includes all norms not derived from an inner product. In the following, the identity map on X is denoted I_X, or simply I.

Lemma 2.13. *Let $\tau : \Gamma \to \operatorname{Isom}(X)$ be a group representation and let $\chi \in \hat{\Gamma}$. If $\mathcal{T}(X)$ is minimal then,*

$$\dim_{\mathbb{C}} \mathcal{T}_\chi(X) = \dim_{\mathbb{C}} \left(\bigcap_{\gamma \in \Gamma} \ker(\tau(\gamma) - \overline{\chi(\gamma)} I) \right).$$

Proof. Let $\eta \in \mathcal{T}(X)$. Since $\mathcal{T}(X)$ is minimal, there exists $x_0 \in X$ such that $\eta(x) = x_0$ for all $x \in X$. Note in particular that $x_0 = \eta(\tau(\gamma)(x_0))$ for each $\gamma \in \Gamma$. Thus $\eta_{\mathbb{C}} \in \mathcal{T}_\chi(X)$ if and only if $x_0 = \chi(\gamma)\tau(\gamma)(x_0)$ for each $\gamma \in \Gamma$. The result now follows. □

Let $\omega = e^{2\pi i/n}$, where $n \in \mathbb{N}$ and $n \geq 2$, and consider the multiplicative cyclic group $\mathbb{Z}_n = \{\omega^k : k = 0, 1, \ldots, n-1\}$. Recall that the dual group for $\mathbb{Z}_n$ consists of characters $\chi_0, \chi_1, \ldots, \chi_{n-1}$ where $\chi_j(\omega) = \omega^j$ for each $j = 0, 1, \ldots, n-1$.

Lemma 2.14. *Let $\tau : \mathbb{Z}_n \to \operatorname{Isom}(X)$ be a group representation where $n \geq 2$. If $\mathcal{T}(X)$ is minimal then, for each $j = 0, 1, \ldots, n-1$,*

$$\dim_{\mathbb{C}} \mathcal{T}_{\chi_j}(X) = \dim_{\mathbb{C}} \ker(\tau(\omega) - \overline{\omega}^j I).$$

Proof. Let $j \in \{0, 1, \ldots, n-1\}$. Note that $\ker(\tau(\omega) - \overline{\omega}^j I) \subseteq \ker(\tau(\omega^k) - \overline{\omega}^{jk} I)$ for $k = 0, 1, \ldots, n-1$. Thus, by Lemma 2.13,

$$\dim_{\mathbb{C}} \mathcal{T}_{\chi_j}(X) = \dim_{\mathbb{C}} \left(\bigcap_{k=0}^{n-1} \ker(\tau(\omega^k) - \overline{\omega}^{jk} I) \right) = \dim_{\mathbb{C}} \ker(\tau(\omega) - \overline{\omega}^j I).$$

□

In the following, an *n-fold rotation* ($n \geq 2$) of a two-dimensional real vector space X is a linear operator $S : X \to X$ with matrix $\begin{pmatrix} \cos(2\pi/n) & -\sin(2\pi/n) \\ \sin(2\pi/n) & \cos(2\pi/n) \end{pmatrix}$ with respect to some basis for X. If $\dim X \geq 3$ then a linear operator $T : X \to X$ is an *n-fold rotation* if there exists a direct sum decomposition $X = Y \oplus Z$, where Y is a two-dimensional subspace of X, with respect to which $T = S \oplus I_Z$, S is an n-fold rotation of Y and I_Z is the identity operator on Z.

Lemma 2.15. *Let $\tau : \mathbb{Z}_n \to \mathrm{Isom}(X)$ be a group representation where $\tau(\omega)$ is an n-fold rotation of X and $n \geq 2$. Suppose, in addition, that $\mathcal{T}(X)$ is minimal.*

(i) If $n = 2$ then,

$$\dim_{\mathbb{C}} \mathcal{T}_{\chi_j}(X) = \begin{cases} \dim_{\mathbb{R}} X - 2 & \text{if } j = 0, \\ 2 & \text{if } j = 1. \end{cases}$$

(ii) If $n \geq 3$ then,

$$\dim_{\mathbb{C}} \mathcal{T}_{\chi_j}(X) = \begin{cases} \dim_{\mathbb{R}} X - 2 & \text{if } j = 0, \\ 1 & \text{if } j \in \{1, n-1\}, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Write $X = Y \oplus Z$ and $\tau(\omega) = S \oplus I_Z$ where $\dim Y = 2$ and S is an n-fold rotation of Y. Then $X_{\mathbb{C}} = Y_{\mathbb{C}} \oplus Z_{\mathbb{C}}$. Note that $\tau(\omega) - \overline{\chi_0(\omega)} I = (S - I_Y) \oplus 0$. Also, $S - I_Y$ is invertible and so, by Lemma 2.14,

$$\dim_{\mathbb{C}} \mathcal{T}_{\chi_0}(X) = \dim_{\mathbb{C}} \ker((S - I_Y) \oplus 0) = \dim_{\mathbb{R}} X - 2.$$

Now let $j \in \{1, \ldots, n-1\}$. If $n = 2$, then $\omega = -1$ and $S = -I_Y$. Note that $\tau(\omega) - \overline{\omega} I = 0 \oplus 2I_Z$ and so, by Lemma 2.14,

$$\dim_{\mathbb{C}} \mathcal{T}_{\chi_1}(X) = \dim_{\mathbb{C}} \ker(0 \oplus 2I_Z) = \dim_{\mathbb{R}} Y = 2.$$

If $n \geq 3$, then S has eigenvalues of multiplicity 1 at ω and $\overline{\omega}$. Note that $\tau(\omega) - \overline{\omega}^j I = (S - \overline{\omega}^j I_Y) \oplus (1 - \overline{\omega}^j) I_Z$ and so, by Lemma 2.14,

$$\dim_{\mathbb{C}} \mathcal{T}_{\chi_j}(X) = \dim_{\mathbb{C}} \ker(\tau(\omega) - \overline{\omega}^j I) = \dim_{\mathbb{C}} \ker(S - \overline{\omega}^j I_Y) = \begin{cases} 1 & \text{if } j \in \{1, n-1\}, \\ 0 & \text{otherwise.} \end{cases}$$

□

2.8. Gain graphs

Let (G, θ) be a Γ-symmetric graph and fix an orientation on the edges of the quotient graph $G_0 = (V_0, E_0)$. For each vertex orbit $[v] \in V_0$, choose a representative vertex $\tilde{v} \in [v]$ and denote the set of all such representatives by $\tilde{V}_0$. For each directed edge $[e] = ([v], [w])$ in the directed multigraph G_0 there exists a unique $\gamma \in \Gamma$, referred to as the *gain* on $[e]$, such that $\tilde{v}(\gamma\tilde{w}) \in [e]$. This gain assignment $\psi : E_0 \to \Gamma$, $[e] \mapsto \psi_{[e]}$, is well-defined and the pair (G_0, ψ) is referred to as a *(quotient) gain graph* for (G, θ). The graph G is also called the *covering graph* of (G_0, ψ).

Note that a gain assignment ψ is dependent on the choice of representative vertices $\tilde{V}_0$ and also on the choice of orientation for each edge of G_0. We may *switch* the gain assignment on the directed multigraph G_0 by choosing a different set of vertex orbit representatives. We regard two gain assignments on the directed multigraph G_0 as equivalent if one can be obtained from the other by such a switching operation. Note that if the orientation of an edge $[e]$ in G_0 is reversed then the induced gain $\psi_{[e]}$ is replaced with $\psi_{[e]}^{-1}$.

In general, we refer to a group-labelled directed multigraph (G_0, ψ) with $\psi : E_0 \to \Gamma$ as a Γ-*gain graph* if it is a quotient gain graph for a Γ-symmetric graph (G, θ). Note that, since G is assumed to be simple, (G_0, ψ) has no parallel edges with the same gain when oriented in the same direction and no loops with a trivial gain. We regard two Γ-gain graphs as equivalent if they are derived from the same Γ-symmetric graph (G, θ). For more on gain graphs we refer the reader to [4, 21].

Example 2.16. Figure 1 illustrates several examples of $\mathbb{Z}_2$-symmetric graphs together with accompanying quotient gain graphs. These gain graphs will form base graphs for the inductive construction presented in Section 3. Note that in the case of $\mathbb{Z}_2$-symmetric graphs, gain assignments are independent of the chosen edge orientation. Thus edge orientations have been omitted from Figure 1. In each gain graph, the indicated gains are determined by the set of representative vertices, labelled by p, in its covering graph. Note that each covering graph is presented as a two-dimensional bar-joint framework with half-turn rotational symmetry. Moreover, it can be shown that these bar-joint frameworks are χ_0-symmetrically isostatic with respect to the ℓ^∞ norm. The reasons for this, and the significance of the edge colourings, are explained in Section 4.1.

The *gain* of a path of directed edges $F = [v_1], [e_1], [v_2], \ldots, [e_k], [v_k]$ in a gain graph (G_0, ψ) (where $[v_1]$ may be equal to $[v_k]$) is defined as the product $\psi(F) = \Pi_{i=1}^{k} \psi([e_i])^{\mathrm{sign}([e_i])}$, where $\mathrm{sign}([e_i]) = 1$ if $[e_i]$ is directed from $[v_i]$ to $[v_{i+1}]$ and $\mathrm{sign}([e_i]) = -1$ if $[e_i]$ is directed from $[v_{i+1}]$ to $[v_i]$. A set of edges F is *balanced* if it does not contain a cycle of edges, or, has the property that every cycle of edges in F has gain 1. A subgraph of G_0 is *balanced* in (G_0, ψ) if its edge set is balanced; otherwise, the subgraph is *unbalanced*.

Lemma 2.17 ([4, 21])**.** *Let G_0 be a quotient graph and fix an orientation on the edges of G_0. If a subgraph H_0 is balanced for some gain assignment on the directed quotient graph G_0 then,*

(i) H_0 is balanced for every equivalent gain assignment on the directed quotient graph G_0.

(ii) there exists an equivalent gain assignment ψ on the directed quotient graph G_0 which satisfies $\psi([e]) = 1$ for all $[e] \in E(H_0)$.

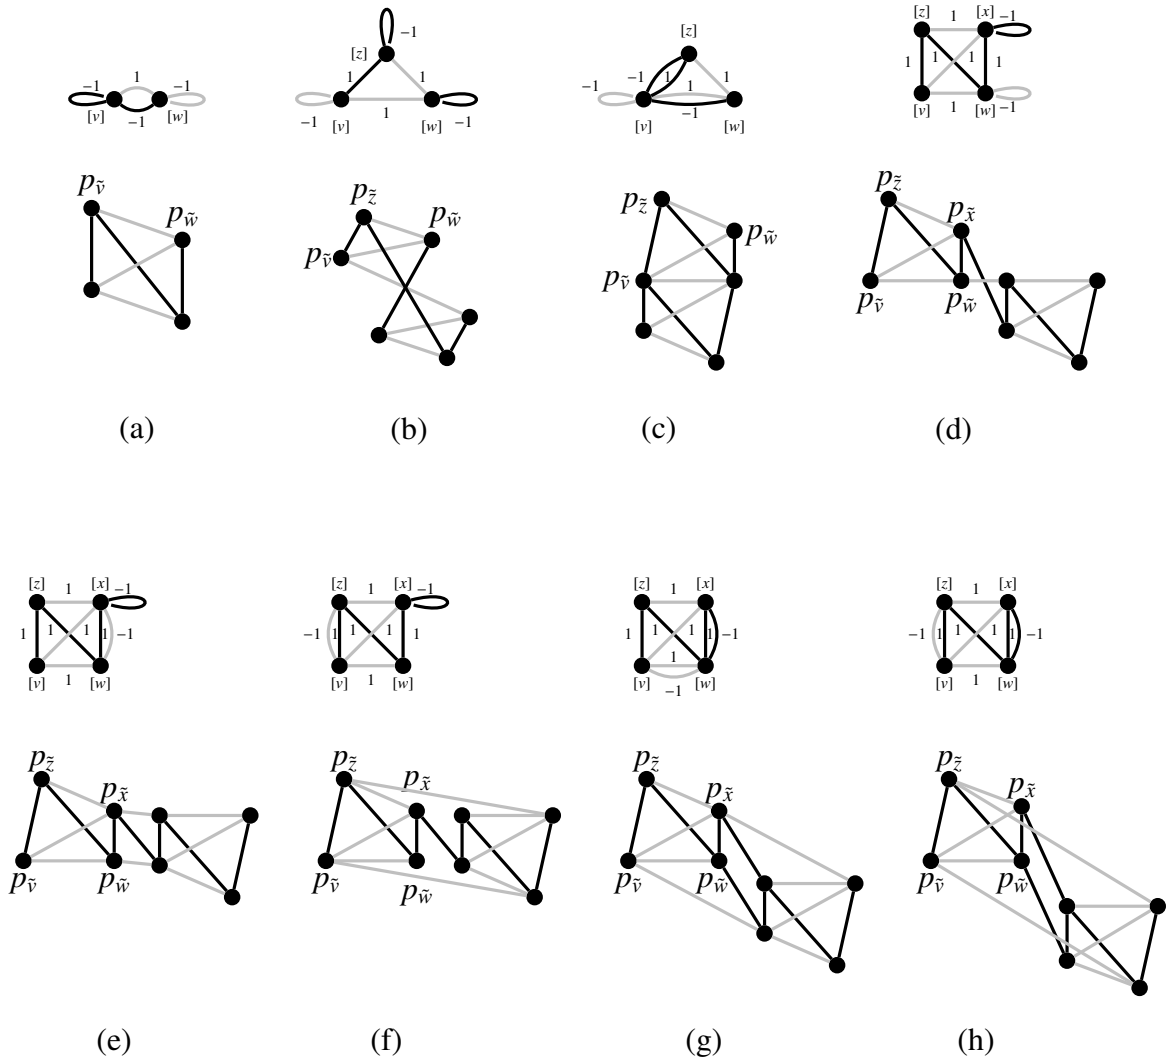

Figure 1: Examples of $\mathbb{Z}_2$-gain graphs and their covering graphs. These are precisely the base $\mathbb{Z}_2$-gain graphs for the $(2,2,0)$-gain-tight inductive construction described in Section 3. The bottom rows illustrate χ_0-symmetrically isostatic realisations for the ℓ^∞-plane under half-turn rotational symmetry. The monochrome subgraphs induced by these realisations (described in Section 4) are indicated in black and grey.

2.9. A special case

Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and Γ-symmetric bar-joint framework in X and suppose the Γ-symmetric graph (G, θ) has an associated gain graph which is balanced. By Lemma 2.17, there exists a choice of vertex orbit representatives $\tilde{V}_0$ such that the induced gain assignment satisfies $\psi([e]) = 1$ for all $[e] \in E(G_0)$. It follows that G_0 is a simple graph. Consider the well-positioned bar-joint framework $(G_0, \tilde{p})$ in X where $\tilde{p}_{[v]} = p_{\tilde{v}}$ for each vertex orbit $[v] \in V_0$ and vertex orbit representative $\tilde{v} \in \tilde{V}_0$ with $\tilde{v} \in [v]$. The following lemma shows the relationship between the differential $df_{G_0}(\tilde{p})$ and the components of the block decomposition of $df_G(p)$ described in Proposition 2.5.

Lemma 2.18. *Let $\chi \in \hat{\Gamma}$ and define a pair of linear transformations,*

$$S_\chi(\mathcal{G}) : (X_{\mathbb{C}})^{V_0} \to X_\chi, \quad (x_{[v]})_{[v]\in V_0} \mapsto (\chi(\gamma)\tau(\gamma)x_{[v]})_{v\in V}$$

where $v = \gamma\tilde{v}$ for some unique $\tilde{v} \in \tilde{V}_0$ and some unique $\gamma \in \Gamma$, and,

$$T_\chi(\mathcal{G}) : \mathbb{C}^{E_0} \to Y_\chi, \quad (x_{[e]})_{[e]\in E_0} \mapsto (\chi(\gamma)x_{[e]})_{e\in E}$$

where $e = \gamma(\tilde{v}\tilde{w})$ for some unique $\tilde{v}, \tilde{w} \in \tilde{V}_0$ and some unique $\gamma \in \Gamma$.

Then the following diagram commutes.

$$\begin{tikzcd}
(X_{\mathbb{C}})^{V_0} \arrow[r, "df_{G_0}(\tilde{p})"] \arrow[d, "S_\chi(\mathcal{G})"'] & \mathbb{C}^{E_0} \arrow[d, "T_\chi(\mathcal{G})"] \\
X_\chi \arrow[r, "R_\chi(\mathcal{G})"] & Y_\chi
\end{tikzcd}$$

In particular, $df_{G_0}(\tilde{p})$ and $R_\chi(\mathcal{G})$ are similar linear transformations.

Proof. Let $u = (u_{[v]})_{[v]\in V_0} \in (X_{\mathbb{C}})^{V_0}$ and let $e = vw \in E$. Then $v = \gamma\tilde{v}$ and $w = \gamma\tilde{w}$ for some unique vertex orbit representatives $\tilde{v}, \tilde{w} \in \tilde{V}_0$ and some unique $\gamma \in \Gamma$. Recall from Proposition 2.5 that $R_\chi(\mathcal{G})$ is the restriction of $df_G(p)$ to the subspace $X_\chi \subset (X_{\mathbb{C}})^V$. Thus, by Lemma 2.1, the e-component of $(R_\chi(\mathcal{G}) \circ S_\chi(\mathcal{G}))(u)$ is given by $\varphi_{v,w}(\chi(\gamma)\tau(\gamma)(u_{[v]} - u_{[w]}))$. Also, by applying Lemma 2.1 to the bar-joint framework $(G_0, \tilde{p})$, and using Lemma 2.4, we see that the e-component of $(T_\chi(\mathcal{G}) \circ df_{G_0}(\tilde{p}))(u)$ is given by $\chi(\gamma)\varphi_{\tilde{v},\tilde{w}}(u_{[v]} - u_{[w]}) = \varphi_{v,w}(\chi(\gamma)\tau(\gamma)(u_{[v]} - u_{[w]}))$. □

2.10. Gain-sparsity

Let $k \in \mathbb{N}$, let $l \in \{0, 1, \ldots, 2k-1\}$ and let $m \in \{0, 1, \ldots, l\}$.

Definition 2.19. A gain graph (G_0, ψ) is (k, l, m)-*gain-sparse* if

(a) $|F| \leq k|V(F)| - l$ for any nonempty balanced $F \subseteq E(G_0)$, and,

(b) $|F| \leq k|V(F)| - m$ for all $F \subseteq E(G_0)$.

Moreover, (G_0, ψ) is (k, l, m)-*gain-tight* if $|E(G_0)| = k|V(G_0)| - m$ and (G_0, ψ) is (k, l, m)-gain-sparse.

Consider again the multiplicative cyclic group $\mathbb{Z}_n = \{\omega^k : k = 0, 1, \ldots, n-1\}$ with characters $\chi_j(\omega) = \omega^j$ for $j = 0, 1, \ldots, n-1$. A $\mathbb{Z}_n$-symmetric bar-joint framework $\mathcal{G} = (G, p, \theta, \tau)$ in X is said to be C_n-*symmetric* if $\tau(\omega)$ is an n-fold rotation of X.

Corollary 2.20. *Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and C_n-symmetric bar-joint framework in X, where $n \geq 2$, and let $d = \dim_{\mathbb{R}} X$. Suppose, in addition, that $\mathcal{T}(X)$ is minimal and $\mathcal{G}$ is χ_j-symmetrically isostatic.*

(i) Suppose $n = 2$.

(a) If $j = 0$ then (G_0, ψ) is $(d, d, d-2)$-gain-tight.

(b) If $j = 1$ then, (G_0, ψ) is $(d, d, 2)$-gain-tight.

(ii) Suppose $n \geq 3$.

(a) If $j = 0$ then (G_0, ψ) is $(d, d, d-2)$-gain-tight.

(b) If $j \in \{1, n-1\}$ then, (G_0, ψ) is $(d, d, 1)$-gain-tight.

(c) If $j \notin \{0, 1, n-1\}$ then, (G_0, ψ) is $(d, d, 0)$-gain-tight.

Proof. Let $\chi = \chi_j$. Note that since $\mathcal{T}(X)$ is minimal, every bar-joint framework in X is full. By Lemma 2.10, $\mathcal{G}$, and every C_n-symmetric subframework of $\mathcal{G}$, is χ-symmetrically full. Also note that $\dim X_\chi = (\dim_{\mathbb{R}} X)|V_0|$ and $\dim Y_\chi = |E_0|$. Let $F \subseteq E(G_0)$, let H_0 be the subgraph of G_0 spanned by the edges in F and let H be the covering graph for H_0 in G. By Lemma 2.11, the C_n-symmetric subframework $\mathcal{H} = (H, p_H, \theta_H, \tau_H)$ is χ-symmetrically independent. Thus, by Proposition 2.12(*ii*),

$$|E(H_0)| \leq (\dim_{\mathbb{R}} X)|V(H_0)| - \dim_{\mathbb{C}} \mathcal{T}_\chi(X).$$

If H_0 is a balanced subgraph of G_0 then we may consider an associated bar-joint framework $(H_0, \tilde{p}_H)$, as described in Section 2.9. By Lemma 2.18, $df_{H_0}(\tilde{p}_H)$ and $R_\chi(\mathcal{H})$ are similar linear transformations. It follows that $(H_0, \tilde{p}_H)$ is an independent subframework of (G, p) and so,

$$\begin{aligned} |E(H_0)| = \operatorname{rank} df_{H_0}(\tilde{p}_H) &= (\dim_{\mathbb{R}} X)|V(H_0)| - \dim_{\mathbb{R}} \mathcal{F}(H_0, \tilde{p}_H) \\ &\leq (\dim_{\mathbb{R}} X)|V(H_0)| - \dim_{\mathbb{R}} \mathcal{T}(X). \end{aligned}$$

Since $\mathcal{T}(X)$ is minimal, $\dim_{\mathbb{R}} \mathcal{T}(X) = \dim_{\mathbb{R}} X$. Thus the results now follow from Lemma 2.15 and Proposition 2.12(*iii*). □

Remark 2.21. Note that, by the above corollary, for two-dimensional χ_j-symmetrically isostatic bar-joint frameworks with rotational symmetry, the associated gain graph must be either $(2, 2, 0)$-gain tight, $(2, 2, 1)$-gain-tight or $(2, 2, 2)$-gain-tight. Inductive constructions for $(2, 2, 1)$- and $(2, 2, 2)$-gain-tight gain graphs are presented in [15] (see also [13, 12]). In the next section we present an inductive construction for $(2, 2, 0)$-gain-tight gain graphs. Also note that, in any dimension, the (k, l, m)-gain tight counts given by Corollary 2.20 are the bases of a matroid as was observed in [15]. (Note however that this matroidal property does not hold for arbitrary triples $k, l, m \in \mathbb{N}$. Indeed it fails in some rigidity contexts [3].)

3. An inductive construction of (2, 2, 0)-gain tight $\mathbb{Z}_2$-gain graphs

Let (G_0, ψ) be a $\mathbb{Z}_2$-gain graph with covering graph G. For simplicity, we will omit the square brackets in the notation of vertices and edges of (G_0, ψ) in this section, and simply write v for the vertex $[v]$, and (uv, α) for the edge $([u], [v])$ with gain α. Note that the orientation of the edges of (G_0, ψ) does not matter, since (G_0, ψ) is a $\mathbb{Z}_2$-gain graph and $\mathbb{Z}_2$ is of order 2. For the remainder of this article we will only consider $\mathbb{Z}_2$-gain graphs and so from now on the term *gain graph* will be used to mean $\mathbb{Z}_2$-*gain graph*.

3.1. Base graphs

Let $\mathcal{B}$ denote the family of $(2, 2, 0)$-gain-tight *base graphs* presented in Figure 1. It will be convenient to assign names to elements of $\mathcal{B}$. Let iK_j^ℓ denote the complete graph on j vertices, with i copies of each edge and ℓ loops on each vertex. Then $2K_2^1$ and K_3^1 with a balanced K_3 are the gain graphs in Figures 1(a) and (b), respectively. The graph formed from $2K_3$ by adding a loop and deleting an edge not incident with the vertex with the loop will be denoted by R. (See Figure 1(c).) We denote by K_4^+ any $\mathbb{Z}_2$-gain graph formed from a balanced copy of K_4 by adding a single edge (subject to $(2, 2, 0)$-gain-sparsity). We shall also use K_4^{++} to denote any one of the non-isomorphic $(2, 2, 0)$-gain-tight gain graphs formed from K_4 by adding two edges. (See Figures 1(d)-(h).)

3.2. Preliminaries

We first record two preliminary lemmas about gain graphs which go back to Zaslavsky [21].

Lemma 3.1. *Let G_0 be a (simple) cycle. A $\mathbb{Z}_2$-gain graph (G_0, ψ) is unbalanced if and only if the vertices in V_0 can be switched so that any one edge has non-identity gain and every other edge in the resulting $\mathbb{Z}_2$-gain graph (G_0, ψ') has identity gain.*

Lemma 3.2. *Let (G_0, ψ) be a $\mathbb{Z}_2$-gain graph and let A and B be subgraphs of (G_0, ψ). Suppose that $A \cap B$ is connected. If A and B are balanced then $A \cup B$ is also balanced.*

We will also need some elementary results about sparse graphs which we record here for convenience. Let $f(G_0) = 2|V_0| - |E_0|$. So, for example, any $(2, 2, 0)$-gain-tight gain graph G_0 satisfies $f(G_0) = 0$ while any balanced subgraph G_0' satisfies $f(G_0') \geq 2$.

Lemma 3.3. *Let G_0 be connected and 4-regular. Then $f(G_0') \geq 1$ for any proper subgraph $G_0' \subset G_0$.*

Proof. Suppose G_0 contains a subgraph G_0' with $f(G_0') = 0$. Then G_0' has average degree 4, so G_0' must be 4-regular (by the 4-regularity of G_0). Since G_0 is connected it follows that $G_0' = G_0$. □

For two disjoint vertex sets $A, B \subset V(G_0)$, we denote by $d(A, B)$ the number of edges between A and B.

Lemma 3.4. *Let $H_0 = (V_0', E_0')$ be a subgraph of G_0. If the degree of v in G_0 is at least 4 for all $v \in V_0'$ then $d(V_0', V_0 - V_0') \geq 2f(H_0)$.*

Proof. Since $|E_0'| = 2|V_0'| - f(H_0)$ and every vertex in V_0' has degree at least 4 in G_0 we have,

$$4|V_0'| \leq \sum_{v \in V_0'} \deg_{G_0}(v) = 2|E_0'| + d(V_0', V_0 - V_0') = 4|V_0'| - 2f(H_0) + d(V_0', V_0 - V_0').$$

□

Lemma 3.5. *Let (G_0, ψ) be $(2,2,0)$-gain-sparse, and let G_0' be a balanced subgraph of (G_0, ψ) with $f(G_0') \in \{2,3\}$. Then G_0' is connected.*

Proof. Suppose G_0' is disconnected. Let A be a connected component of G_0' and let $B = G_0' - A$. Since any subgraph of a balanced gain graph is also balanced, we have $f(A) \geq 2$ and $f(B) \geq 2$. Hence $f(G_0') = f(A) + f(B) \geq 4$, contradicting the hypothesis of the lemma. □

Lemma 3.6. *Let (G_0, ψ) be $(2,2,0)$-gain-tight. Let H_0' and H_0'' be balanced subgraphs of (G_0, ψ) with $V(H_0') \cap V(H_0'') \neq \emptyset$ and $f(H_0') = 2 = f(H_0'')$. Then either*

(i) $f(H_0' \cap H_0'') = 4$ *and* $f(H_0' \cup H_0'') = 0$*, or,*

(ii) $f(H_0' \cup H_0'') = f(H_0' \cap H_0'') = 2$.

Moreover, (ii) holds if and only if $H_0' \cup H_0''$ is balanced.

Proof. As $H_0' \cap H_0'' \subset H_0'$ we have $f(H_0' \cap H_0'') \geq 2$. If $f(H_0' \cap H_0'') \geq 4$ then

$$0 \leq f(H_0' \cup H_0'') = f(H_0') + f(H_0'') - f(H_0' \cap H_0'') = 4 - f(H_0' \cap H_0'') \leq 0$$

and so (*i*) holds. If $f(H_0' \cap H_0'') \in \{2,3\}$ then $H_0' \cap H_0''$ is connected by Lemma 3.5. It follows that $H_0' \cup H_0''$ is balanced by Lemma 3.2 and hence $f(H_0' \cup H_0'') \geq 2$. Thus,

$$2 \leq f(H_0' \cup H_0'') = 2 + 2 - f(H_0' \cap H_0'') \leq 2$$

and so (*ii*) holds. □

3.3. Henneberg-type operations

Now we define operations on $\mathbb{Z}_2$-gain graphs. The H1 operation (or Henneberg 1 move, or 0-extension) adds a new vertex of 3 possible types. In type 1a the new vertex has degree 2 and two distinct neighbours; in type 1b the new vertex has degree 2 and one neighbour with two parallel edges; and in type 1c the new vertex has degree 3 with one neighbour and a loop. (See Figure 2.) The gains on the new edges are arbitrary subject to the condition that the covering graph is simple, i.e. parallel edges have different gains and a loop has gain -1.

The H2 operation (or Henneberg 2 move, or 1-extension) deletes one edge (xy, α) and adds a new vertex v adjacent to x, y of five possible types. In type 2a, v has degree 3 and 3 distinct neighbours with edges (xv, β) and (yv, γ) satisfying $\beta\gamma = \alpha$; in type 2b, v has degree 3 and exactly 2 neighbours with edges $(xv, 1), (xv, -1)$ and (yv, δ) with $\delta = \pm 1$; in type 2c, the deleted edge xy is a loop $(xx, -1)$ and v has degree 3 and exactly 2 neighbours with edges $(xv, 1), (xv, -1)$ and (yv, δ) with $\delta = \pm 1$; in type 2d, v has degree 4 and exactly 2 neighbours with edges $(xv, \beta), (yv, \gamma)$ and

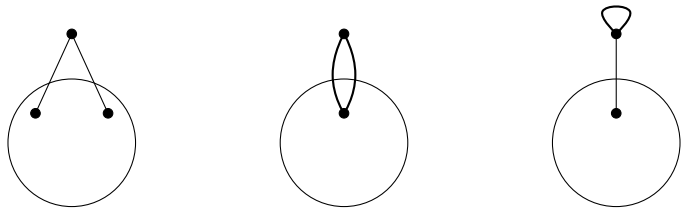

Figure 2: H1 a, b, c operations on gain graphs. Gain labels are omitted.

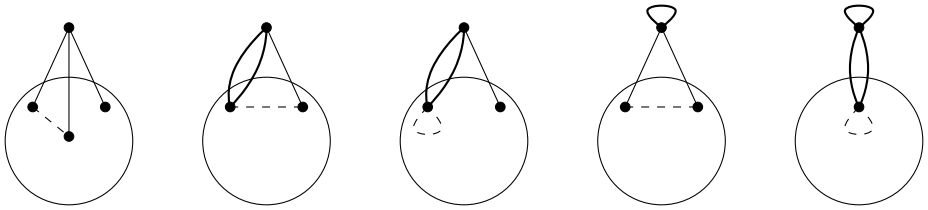

Figure 3: H2 a, b, c, d, e operations on gain graphs. Gain labels are omitted.

$(vv, -1)$ satisfying $\alpha = \beta\gamma$; finally, in type 2e, the deleted edge xy is a loop $(xx, -1)$, v has degree 4 and exactly 1 neighbour with edges $(xv, 1), (xv, -1)$ and $(vv, -1)$. (See Figure 3.)

The H3 operation (or X-replacement, or 2-extension) deletes two edges $(xy, \alpha), (zw, \beta)$ and adds a new degree 4 vertex v adjacent to x, y, z, w of five possible types. In type 3a, v has 4 distinct neighbours and edges $(xv, \gamma), (yv, \delta), (zv, \epsilon), (wv, \zeta)$ where $\alpha = \gamma\delta$ and $\beta = \epsilon\zeta$; in type 3b, v has 3 distinct neighbours, $y = z$ and there are two parallel edges between v and y, with edges $(xv, \gamma), (yv, 1), (yv, -1), (wv, \zeta)$ where $\alpha = \gamma$ and $\beta = -\zeta$; in type 3c, v has 3 distinct neighbours, $x = y$ so $\alpha = -1$ and there are two parallel edges between v and x with edges $(xv, -1), (xv, 1), (zv, \epsilon), (wv, \zeta)$ and $\beta = \epsilon\zeta$; in type 3d, v has 2 distinct neighbours, $x = y$ and $z = w$ so $\alpha = \beta = -1$ and there are two parallel edges between v and x and between v and z with edges $(xv, 1), (xv, -1), (zv, 1), (zv, -1)$. (See Figure 4.)

Figure 4: H3 a, b, c, d operations on gain graphs. Gain labels are omitted.

A *vertex-to-K_4* operation removes a vertex v (of arbitrary degree) and all the edges incident with v, and adds in a copy of K_4 with only trivial gains. Each removed edge (xv, γ), where $x \neq v$, is replaced by an edge (xy, γ) for some y in the new K_4. If $x = v$ then the removed edge $(vv, -1)$ is replaced by an edge $(wz, -1)$ where w, z are vertices of the new K_4. Note w and z need not be distinct. (See Figure 5.)

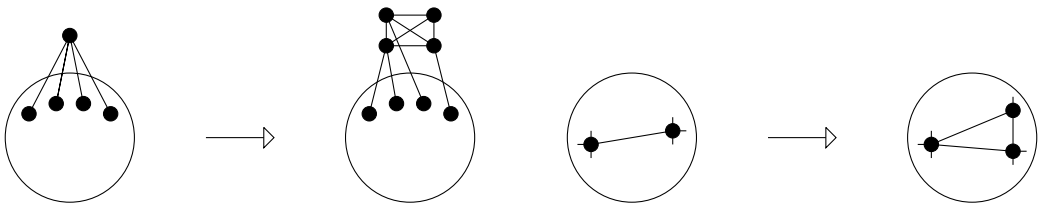

Figure 5: The vertex-to-K_4 operation and the vertex splitting operation. Gain labels are omitted.

A *vertex splitting* operation first chooses a vertex v_1, a neighbour v_2 of v_1, and a partition N_1, N_2 of the remaining neighbours of v_1; it then deletes the edges from v_1 to vertices in N_1, adds a new vertex v_0 joined to vertices in N_1 and finally adds two new edges v_0v_1, v_0v_2. If there is a loop at v_1 then it is either left unchanged or replaced by a loop at v_0. We specify that v_0v_1 is given gain 1 and v_0v_2 is given the same gain as v_1v_2. (See Figure 5.)

By construction we have the following.

Lemma 3.7. *Applying any of the above operations to a* $(2,2,0)$*-gain-tight gain graph results in a* $(2,2,0)$*-gain-tight gain graph.*

Proof. When the operation is a H1, H2 or H3 operation then we may employ similar arguments to those in [4, Lemma 4.1 and 7.6].

Suppose (G,ψ) is obtained from (H,ψ') by a vertex-to-K_4 operation at v. If (G,ψ) is not $(2,2,0)$-gain-tight then there exists a vertex-induced subgraph G_1 of (G,ψ) such that $f(G_1) < 0$, or, $(G_1, \psi|_{G_1})$ is balanced and $f(G_1) \in \{0,1\}$. Consider the subgraph H_1 of H corresponding to G_1 obtained on contracting G to H. Note that in our definition of H_1, if there is a loop at v in H then H_1 will contain this loop if and only if G_1 contains the extra edge $(wz, -1)$ in the copy of K_4^+. Note that $(H_1, \psi'|_{H_1})$ is balanced if and only if $(G_1, \psi|_{G_1})$ is balanced. There are two possibilities: either $|V(K_4 \cap G_1)| \in \{1,4\}$, or, $|V(K_4 \cap G_1)| \in \{2,3\}$. In the first case, $f(H_1) = f(G_1)$ and in the second case $f(H_1) = f(G_1) - 1$, contradicting $(2,2,0)$-gain-sparsity.

Lastly, suppose (G,ψ) is obtained from a $(2,2,0)$-gain-tight gain graph (H,ψ') by a vertex splitting operation at the vertex v_1 which adjoins the new vertex v_0. Suppose that (G,ψ) is not $(2,2,0)$-gain-tight. Since $f(G) = 0$ it follows that there exists a vertex-induced subgraph G_1 of (G,ψ) such that $f(G_1) < 0$, or, $(G_1, \psi|_{G_1})$ is balanced and $f(G_1) \in \{0,1\}$. Consider the subgraph H_1 of H corresponding to G_1 obtained on contracting G to H. Note that if G_1 contains v_0 and v_1, then $f(H_1) \leq f(G_1)$. Otherwise, $f(H_1) = f(G_1)$. Hence H_1 violates $(2,2,0)$-gain-sparsity. □

3.4. Reducing low-degree vertices via reverse Henneberg-type operations

Note that if a gain graph G_0' is obtained from a $(2,2,0)$-gain-tight gain graph by reversing any of the above operations then $f(G_0') = 0$. Thus G_0' is $(2,2,0)$-gain-tight if and only if each subgraph of G_0' satisfies the $(2,2,0)$-sparsity counts. A vertex v in a $(2,2,0)$-gain-tight gain graph is *admissible* if there is a reverse $H1$ operation, a reverse $H2$ operation or a reverse $H3$ operation removing v which results in a $(2,2,0)$-gain-tight gain graph. Similarly a balanced subgraph isomorphic to K_4 or K_3 is *admissible* if there is a K_4-contraction (i.e. a reverse vertex-to-K_4 operation) or edge

contraction (i.e. a reverse vertex splitting operation) which results in a (2,2,0)-gain-tight gain graph.

Our first lemma is trivial and deals with all H1 moves.

Lemma 3.8. *Let (G_0, ψ) be a $(2, 2, 0)$-gain-tight gain graph. Suppose $v \in V_0$ has degree 2 or is incident to a loop and has degree 3. Then v is admissible.*

We now work through the H2 moves in turn.

Lemma 3.9. *Let (G_0, ψ) be a $(2, 2, 0)$-gain-tight gain graph. Suppose $v \in V_0$ has degree 3 with exactly three neighbours a, b, c. Then v is admissible if and only if it is not contained in a balanced subgraph isomorphic to K_4.*

Proof. Suppose v is admissible. Then there exists a $(2, 2, 0)$-gain-tight gain graph (G'_0, ψ') which is the result of a reverse H2a operation at v. If v is contained in a balanced K_4 subgraph then the deleted edge in (G'_0, ψ') must be one of two parallel edges with equal gain, contradicting the simplicity of the covering graph for (G'_0, ψ').

For the converse, suppose v is not contained in a balanced subgraph isomorphic to K_4. Then there exists a gain graph (G'_0, ψ') which is the result of a reverse H2a operation at v. Let (av, α), (bv, β) and (cv, γ) be in E_0.

We first show that if there exists a subgraph H_{ab} of $G_0 - v$ which contains a, b with $f(H_{ab}) = 0$, then v is admissible. Suppose for a contradiction that such a subgraph H_{ab} of G_0 exists and that v is not admissible. If $c \in V(H_{ab})$ then $f(H_{ab} \cup v) < 0$, which contradicts $(2, 0)$-sparsity. So $c \notin V(H_{ab})$. If the edges $(ac, \alpha\gamma), (bc, \beta\gamma)$ are in G_0 then the union of H_{ab} with v, c and the edges ac, bc, va, vb, vc violates $(2, 0)$-sparsity. So we may suppose $(ac, \alpha\gamma) \notin E(G_0)$. Then, since v is not admissible, there exists a subgraph H_{ac} of $G_0 - v$ containing a, c such that either H_{ac} is balanced (with every path from a to c in H_{ac} having gain $\alpha\gamma$) and $f(H_{ac}) = 2$, or $f(H_{ac}) = 0$. In both cases, we clearly have $f(H_{ab} \cup H_{ac} \cup v) < 0$, which is a contradiction.

A similar argument holds for the pairs a, c and b, c. Thus we may assume for each pair of vertices $s, t \in \{a, b, c\}$ that there is no subgraph H_{st} of $G_0 - v$ which contains s, t and satisfies $f(H_{st}) = 0$.

Now assume, without loss of generality, that $(ab, \alpha\beta) \notin E_0$. Suppose there does not exist a balanced subgraph H_{ab} of $G_0 - v$ which contains a, b with $f(H_{ab}) = 2$ and all paths from a to b having gain $\alpha\beta$. Then v is admissible since adding the edge $(ab, \alpha\beta)$ will not violate $(2, 2, 0)$-gain-sparsity.

Suppose $G_0 - v$ does contain a balanced subgraph H_{ab} which contains a, b with $f(H_{ab}) = 2$ and all paths from a to b having gain $\alpha\beta$. We may assume by gain switching (Lemma 2.17) that all edges of H_{ab} have gain 1. In this case, all paths from a to b have gain 1. If $\alpha \neq \beta$ then we claim that v is admissible. To see this note that if there exists a balanced subgraph H'_{ab} which contains a, b with $f(H'_{ab}) = 2$ and all paths from a to b having gain -1 then $H_{ab} \cap H'_{ab}$ is not connected. Thus, by Lemma 3.6, $f(H_{ab} \cup H'_{ab}) = 0$, contradicting our assumption above. Thus v is admissible since we may add the edge $(ab, -1)$. We may now assume $\alpha = \beta$. We may further assume, by gain switching at v, that $\alpha = \beta = 1$.

If H_{ab} contained c then $f(H_{ab} \cup v) = 1$. If vc has gain 1 then we contradict $(2, 2, 0)$-gain-sparsity and so $\gamma = -1$. If $(ac, -1)$ and $(bc, -1)$ are both in (G_0, ψ) then the induced subgraph on

$V(H_{ab}) \cup \{v\}$ violates $(2,0)$-sparsity. Without loss of generality, suppose $(ac,-1)$ is not in (G_0,ψ). Then v is admissible unless there exists a balanced subgraph H_{ac} of $G_0 - v$ which contains a,c with $f(H_{ac}) = 2$ and all paths from a to c having gain -1.

Suppose $G_0 - v$ does contain a balanced subgraph H_{ac} which contains a,c with $f(H_{ac}) = 2$ and all paths from a to c having gain -1. Note that $f(H_{ab} \cup H_{ac}) \geq 1$ (for otherwise adding v and its three edges would contradict $(2,2,0)$-gain sparsity). Thus, by Lemma 3.6, $H_{ab} \cup H_{ac}$ is balanced. By Lemma 3.5, H_{ab} and H_{ac} are connected and so it follows that $H_{ab} \cup H_{ac}$ contains an unbalanced cycle. This is a contradiction and so v is admissible since adding the edge $(ac,-1)$ will not violate $(2,2,0)$-gain-sparsity.

Now suppose H_{ab} does not contain c. In this case, by gain switching at the vertex c (Lemma 2.17) we can assume $\gamma = 1$. If $(ac,1)$ and $(bc,1)$ are both in G_0 then the union of H_{ab} with vertices v,c and the five edges $(ac,1),(bc,1),(va,1),(vb,1),(vc,1)$ is balanced and violates $(2,2)$-sparsity. This is a contradiction and so, without loss of generality, we may assume $(ac,1) \notin G_0$. Then v is admissible unless there exists a balanced subgraph H_{ac} of $G_0 - v$ which contains a,c with $f(H_{ac}) = 2$ and all paths from a to c having gain 1.

Suppose there exists a balanced subgraph H_{ac} of $G_0 - v$ containing a,c with $f(H_{ac}) = 2$, with all paths in H_{ac} from a to c having gain 1. By the previous argument we may assume b is not in H_{ac} since otherwise v is admissible. Since $f(H_{ab} \cup H_{ac}) \geq 1$ (for otherwise adding v and its three edges would contradict $(2,2,0)$-gain sparsity), Lemma 3.6 may be applied, that is $H_{ab} \cup H_{ac}$ is balanced. By gain switching, we may assume the edges of $H_{ab} \cup H_{ac}$ all have gain 1. Let α',β',γ' be the resulting gains on the edges va, vb, vc respectively. If $\alpha' \neq \beta'$ then, by the above argument, v is admissible and we may add the edge $(ab,-1)$. Similarly, v is admissible if $\alpha' \neq \gamma'$. So now suppose $\alpha' = \beta' = \gamma'$. Then $H_{ab} \cup H_{ac} \cup v$ is balanced with $f(H_{ab} \cup H_{ac} \cup v) = 1$. This contradicts $(2,2,0)$-sparsity. □

Lemma 3.10. *Let (G_0,ψ) be a $(2,2,0)$-gain-tight gain graph. Suppose $v \in V_0$ has degree 3 with exactly two neighbours a,b. Then v is admissible if and only if it is not contained in a subgraph isomorphic to R (recall Fig. 1(c)).*

Proof. If v is contained in a subgraph isomorphic to R, then v is clearly not admissible. For the converse, suppose that v is not in a subgraph isomorphic to R. Then there exists a gain graph (G_0',ψ') which is the result of either a reverse H2b operation at v or a reverse H2c operation at v. Let $(av,1),(av,-1)$ and (bv,α) be in E_0.

We first consider a reduction move at v which adds an edge between a and b. Observe that any subgraph H of $G_0 - v$ containing a and b has $f(H) > 0$ (otherwise $f(H \cup v) < 0$ would hold) so we need only consider balanced subgraphs.

Suppose there is no edge ab. If there is no admissible reverse H2b move then there exist distinct balanced subgraphs H_1, H_2 of $G_0 - v$ such that $a,b \in V(H_i)$, $f(H_i) = 2$ for $i = 1,2$ and all paths in H_i from a to b have gain $(-1)^i$. Since $f(H_1 \cap H_2) \geq 2$ either $H_1 \cap H_2$ is not connected and Lemma 3.6 implies that $f(H_1 \cap H_2) = 0$ and adding v violates $(2,0)$-sparsity or $H_1 \cap H_2$ is connected. Then Lemma 3.6 implies that $H_1 \cup H_2$ is balanced. Hence all paths from a to b in $H_1 \cap H_2$ have two distinct gains, a contradiction. Thus v is admissible.

Secondly, suppose there is exactly one edge (ab,β) in E_0. If there is no admissible reverse H2b move, then there exists a balanced subgraph H_{ab} of (G_0,ψ) containing a,b but not v with

$f(H_{ab}) = 2$ such that all paths in H_{ab} from a to b have gain $-\beta$. Note that a does not have a loop (otherwise $(2,0)$-sparsity would be violated). Also, a is not contained in a subgraph H with $f(H) = 0$ (otherwise adjoining the edge (ab, β) to $H \cup H_{ab}$ will violate $(2,0)$-sparsity). Thus a reverse H2c move can be applied which preserves $(2,2,0)$-gain-sparsity. Thus v is again admissible.

Finally, if both $(ab, 1)$ and $(ab, -1)$ are in E_0 then the reverse H2c move adding the loop $(aa, -1)$ is non-admissible if and only if there is a subgraph H of $G - v$ containing a which has $f(H) = 0$. Note that H does not contain b and so $f(H \cup \{v, b\}) = f(H) - 1$, giving a contradiction. □

Lemma 3.11. *Let (G_0, ψ) be a $(2,2,0)$-gain-tight gain graph. Suppose $v \in V_0$ has degree 4 with exactly one loop at v and one neighbour a. Then v is admissible if and only if v is not contained in a subgraph isomorphic to $2K_2^1$.*

Proof. If v is not contained in a subgraph isomorphic to $2K_2^1$ then it is easy to check that v is admissible for a reverse H2e move adding a loop on a. Conversely, if v is contained in a subgraph isomorphic to $2K_2^1$ then G_0 contains a loop at a and so v is clearly not admissible. □

We now move on to H3 moves. First consider the reverse H3d move.

Lemma 3.12. *Let (G_0, ψ) be a $(2,2,0)$-gain-tight gain graph which is 4-regular. Suppose $v \in V_0$ has no loop and exactly two neighbours a, b with a double edge to each. Then v is admissible if and only if G_0 does not contain a loop at a and does not contain a loop at b.*

Proof. Clearly, if G_0 contains a loop at a or b then a reverse H3d move cannot be applied and so v is not admissible. For the converse, suppose G_0 does not contain any loops at a and b. If adding loops on a and b violates $(2,2,0)$-gain-sparsity then either there exists a subgraph H_0 of $G_0 - v$ containing a and b with $f(H_0) = 1$, or there exists a subgraph H_0 of $G_0 - v$ containing a (or b) with $f(H_0) = 0$. In both cases we may use 4-regularity to get a contradiction. If $f(H_0) = 0$ then a must have degree 2 in H_0. Note that H_0 has average degree 4, giving a vertex $c \in H_0$ with degree greater than 4 in H_0, and hence in G_0. This is a contradiction. If H_0 contains a and b and $f(H_0) = 1$, then a and b both have degree 2 in H_0. All other vertices in H_0 have degree at most 4. In this case, since $|E(H_0)| = 2|V(H_0)| - 1$, we obtain the contradiction,

$$\sum_{v \in V(H_0)} \deg_{H_0}(v) \leq 4(|V(H_0)| - 2) + 2 + 2 < 4|V(H_0)| - 2 = 2|E(H_0)|.$$

□

Lemma 3.13. *Let (G_0, ψ) be a $(2,2,0)$-gain-tight gain graph which is connected and 4-regular. Suppose $v \in V_0$ has exactly three neighbours a, b, c and no loop. Suppose no neighbour of v is admissible, then v is admissible if and only if v is not contained in a subgraph isomorphic to K_4^+.*

Proof. If v is contained in a subgraph isomorphic to K_4^+, then v is clearly not admissible. For the converse, suppose v is not contained in a subgraph isomorphic to K_4^+. Let $(av, 1)$, $(av, -1)$, (bv, β), (cv, γ) be the edges incident to v. By Lemma 3.3, (G_0, ψ) does not contain any proper subgraph H with $f(H) = 0$. Note also that no subgraph H of $G_0 - v$ containing $N(v) = \{a, b, c\}$ can have $f(H) = 1$. We may suppose there is no loop on a (otherwise we can use Lemma 3.11 to see that

a is admissible, which is a contradiction). Moreover, we may assume that one of the edges $(ac, 1)$ and $(ac, -1)$ is not in G_0 (otherwise Lemma 3.12 would imply that a is admissible). Similarly, we may assume that one of the edges $(ab, 1)$ and $(ab, -1)$ is not in G_0.

First observe that there are three possible reduction moves: a reverse H3c move which adds a loop at a and the edge $(bc, \beta\gamma)$, a reverse H3b move which adds the edges $(ab, \beta), (ac, -\gamma)$ and a reverse H3b move which adds the edges $(ab, -\beta), (ac, \gamma)$. The reverse H3c move is admissible unless:
(a) the edge $(bc, \beta\gamma)$ already exists in G_0; or
(b) there is a balanced subgraph H_{bc} of $G_0 - v$ containing b, c with $f(H_{bc}) = 2$ in which all paths from b to c have gain $\beta\gamma$.

The reverse H3b move which adds the edges $(ab, \beta), (ac, -\gamma)$ is admissible unless:
(c) one of the edges $(ab, \beta), (ac, -\gamma)$ already exists in G_0; or
(d) there is a balanced subgraph H_{ab} of $G_0 - v$ containing a, b with $f(H_{ab}) = 2$ in which all paths from a to b have gain β, or, there is a balanced subgraph H_{ac} of $G_0 - v$ containing a, c with $f(H_{ac}) = 2$ in which all paths from a to c have gain $-\gamma$.

The reverse H3b move which adds the edges $(ab, -\beta), (ac, \gamma)$ is admissible unless:
(e) one of the edges $(ab, -\beta), (ac, \gamma)$ already exists in G_0; or
(f) there is a balanced subgraph H_{ab} of $G_0 - v$ containing a, b with $f(H_{ab}) = 2$ in which all paths from a to b have gain $-\beta$, or, there is a balanced subgraph H_{ac} of $G_0 - v$ containing a, c with $f(H_{ac}) = 2$ in which all paths from a to c have gain γ.

Suppose (a) holds. Note that if (ab, β) and (ac, γ) both exist or $(ab, -\beta)$ and $(ac, -\gamma)$ both exist then there is a K_4^+ containing v, which is a contradiction. Thus we may suppose that either (ab, β) and $(ac, -\gamma)$ do not exist, or, $(ab, -\beta)$ and (ac, γ) do not exist.

Without loss of generality, we suppose that (ab, β) and $(ac, -\gamma)$ do not exist. If v is not admissible then there exists either a balanced subgraph H_{ab} of $G_0 - v$ containing a, b with $f(H_{ab}) = 2$ and all paths from a to b have gain β or a balanced subgraph H_{ac} of $G_0 - v$ containing a, c with $f(H_{ac}) = 2$ and all paths from a to c having gain $-\gamma$. If both exist then $H = H_{ab} \cup H_{ac}$ is a proper subgraph of G_0 and hence, by Lemma 3.3, $f(H) \geq 1$. Thus, by Lemma 3.6, H is balanced and $f(H) = 2$. Note that $(bc, \beta\gamma) \in E(H)$ (otherwise adjoining this edge to $H \cup v$ violates $(2, 0)$-sparsity). By switching, we may assume without loss of generality that all gains on H are 1. In this case note that $\beta = \gamma$. Adjoining v and the edges $(va, \beta), (vb, \beta), (vc, \beta)$ to H results in a balanced subgraph which violates $(2, 2)$-sparsity. This is a contradiction and so either H_{ab} or H_{ac} does not exist. Without loss of generality we assume H_{ab} does not exist. If H_{ac} does not exist then v is admissible. So suppose H_{ac} does exist.

Suppose $(ab, -\beta) \in E(G_0)$. Then H_{ac} contains the vertex b. (If not, then the vertex a would have degree at most 1 in H_{ac}, a contradiction.) By switching, we may assume all gains in H_{ac} are 1. Moreover, $(bc, \beta\gamma) \in E(H_{ac})$ for otherwise $(2, 0)$-sparsity is violated. Thus $\beta = \gamma$. Now adjoining v together with the edges $(va, \beta), (vb, \beta), (vc, \beta)$ to H_{ac} results in a balanced subgraph which violates $(2, 2)$-sparsity. We conclude that $(ab, -\beta) \notin E(G_0)$. If $(ac, \gamma) \in E(G_0)$ then $(ac, \gamma) \notin E(H_{ac})$, and hence a has degree at most 1 in H_{ac}, a contradiction. Thus $(ac, \gamma) \notin E(G_0)$.

If v is not admissible then there exists either a balanced subgraph H'_{ab} of $G_0 - v$ containing a, b with $f(H'_{ab}) = 2$ and all paths from a to b have gain $-\beta$ or a balanced subgraph H'_{ac} of $G_0 - v$ containing a, c with $f(H'_{ac}) = 2$ and all paths from a to c having gain γ. If H'_{ab} exists

then $H = H'_{ab} \cup H_{ac}$ is a proper subgraph of G_0 and hence, by Lemma 3.3, $f(H) \geq 1$. Thus, by Lemma 3.6, H is balanced and $f(H) = 2$. Note that the edge $(bc, \beta\gamma)$ is not in H, for otherwise we may assume, by switching, that all gains in H are 1, and hence $\beta = \gamma$. As above, this contradicts the $(2,2)$-sparsity of balanced subgraphs of G_0. Now, adjoining the edge $(bc, \beta\gamma)$ to H results in a subgraph of $G_0 - v$ with $f(H) = 1$ which contains $\{a, b, c\}$. This is a contradiction. If H'_{ac} exists then $H = H_{ac} \cup H'_{ac}$ is balanced. However, H contains an unbalanced cycle obtained by concatenating a path from a to c in H_{ac} and a path from a to c in H'_{ac}. This is again a contradiction. We conclude that v is admissible since we may apply a reverse H3b move.

Now suppose (a) does not hold. Consider a reduction at v adding a loop at a and the edge $(bc, \beta\gamma)$. If this move is not admissible then there is a balanced subgraph H_{bc} of $G_0 - v$ containing b, c with $f(H_{bc}) = 2$ in which all paths from b to c have gain $\beta\gamma$. By switching, we may assume without loss of generality that all gains of H_{bc} equal 1. In this case $\beta = \gamma$ since all paths in H_{bc} from b to c have gain 1. We may assume that H_{bc} does not contain a (otherwise adjoining v to H_{bc} results in a balanced subgraph of (G_0, ψ) which violates $(2,2)$-sparsity).

Suppose $ab, ac \in E(G_0)$ with arbitrary gains. The subgraph H obtained from H_{bc} by adjoining the vertices a, v together with the edges $(av, 1), (av, -1)$ and all edges between $\{a, v\}$ and H_{bc} has $f(H) = 0$. Hence $H = G_0$ and the only edges between H_{bc} and a are the edges ab and ac. If ab and ac have the same gain, say δ, then note that $H' = H - (va, -\beta\delta)$ is balanced and satisfies $f(H') = 1$. This is a contradiction. If ab and ac have different gains then, without loss of generality, we can say $(ab, 1)$ and $(ac, -1)$ are in G_0. Since $H = G_0$ and all gains on H_{bc} are 1, there is no path in $G_0 - a$ between b and c with gain -1. It follows that there is an admissible reduction at a adding a loop at v and the edge $(bc, -1)$. This is a contradiction since no neighbour of v is admissible.

Without loss of generality, we may now assume there is no edge between a and b in G_0. Consider the reduction at v adding $(ab, 1)$ and $(ac, -1)$. Suppose there is a balanced subgraph H_{ab} of $G_0 - v$ containing a, b with $f(H_{ab}) = 2$. Note that $c \notin V(H_{ab})$ since otherwise adjoining v to H_{ab} results in a balanced subgraph which violates $(2,2)$-sparsity. The subgraph H obtained from the union of H_{ab} with H_{bc} and the edges $(va, \beta), (vb, \beta), (vc, \beta)$ is balanced and violates $(2,2)$-sparsity. This is a contradiction and so H_{ab} does not exist. Similarly, there is no balanced subgraph H_{ac} of $G_0 - v$ containing a, c with $f(H_{ac}) = 2$. Hence if the reduction at v adding $(ab, 1)$ and $(ac, -1)$ is not admissible then $(ac, -1) \in E(G_0)$. By a similar argument, if the reduction at v adding $(ab, -1)$ and $(ac, 1)$ is not admissible then $(ac, 1) \in E(G_0)$. This is a contradiction since G_0 does not contain both $(ac, 1)$ and $(ac, -1)$, for otherwise a would be admissible. □

Lemma 3.14. *Let (G_0, ψ) be a $(2,2,0)$-gain-tight gain graph. Suppose $v \in V_0$ has degree 4 with exactly four neighbours a, b, c, d. Let $(va, \alpha), (vb, \beta), (vc, \gamma)$ and (vd, δ) be the edges incident to v. Then a reduction at v adding $(ab, \alpha\beta), (cd, \gamma\delta)$ is non-admissible if and only if one of the following conditions holds:*

(i) there is a subgraph H_{ab} of G_0 containing a, b but not v with $f(H_{ab}) = 0$ or there is a subgraph H_{cd} of G_0 containing c, d but not v with $f(H_{cd}) = 0$;

(ii) there is a balanced subgraph H_{ab} of (G_0, ψ) containing a, b but not v with $f(H_{ab}) = 2$ and every path from a to b in H_{ab} has gain $\alpha\beta$ or there is a balanced subgraph H_{cd} of (G_0, ψ) containing c, d but not v with $f(H_{cd}) = 2$ and every path from c to d in H_{cd} has gain $\gamma\delta$;

(iii) there is a balanced subgraph H of (G_0, ψ) containing $N(v)$ but not v with $f(H) = 3$ where every path from a to b in H has gain $\alpha\beta$ and every path from c to d in H has gain $\gamma\delta$;

(iv) one of the edges $(ab, \alpha\beta), (cd, \gamma\delta)$ already exists in (G_0, ψ).

Proof. If any one of these conditions holds then it is clear that v is not admissible. Conversely, suppose that the result of a reduction at v adding $(ab, \alpha\beta), (cd, \gamma\delta)$ is the gain graph (G_0', ψ'). If the reduction is non-admissible then (G_0', ψ') is not $(2, 2, 0)$-gain-tight. It follows that either the covering graph of (G_0', ψ') is not simple, there is a subgraph which violates $(2, 0)$-sparsity or there is a subgraph which violates balanced $(2, 2)$-sparsity. In the first case (iv) holds. In the second case, suppose H' is a subgraph of (G_0', ψ') which violates $(2, 0)$-sparsity. Either H' received one edge in the reduction move and (i) holds or H' received two edges and there is a subgraph H of (G_0, ψ) containing $N(v)$ but not v with $f(H) \leq 1$. However $f(H \cup v) < 0$, contradicting $(2, 2, 0)$-gain-sparsity.

Hence it remains to consider the last case. There is a subgraph H' of (G_0', ψ') which violates balanced $(2, 2)$-sparsity. Therefore H' contains $(ab, \alpha\beta)$ or $(cd, \gamma\delta)$ or both. By symmetry we may suppose $(ab, \alpha\beta) \in H'$. If H' does not contain $(cd, \gamma\delta)$ then this implies that there is a balanced subgraph H of (G_0, ψ) containing a, b but not v with $f(H) = 2$ and every path from a to b in H has gain $\alpha\beta$, giving (ii). Similarly if H' contains $(ab, \alpha\beta)$ and $(cd, \gamma\delta)$ then there is a balanced subgraph H of (G_0, ψ) containing $N(v)$ but not v with $f(H) = 3$ where every path from a to b in H has gain $\alpha\beta$ and every path from c to d in H has gain $\gamma\delta$, giving (iii). □

Lemma 3.15. *Let (G_0, ψ) be a $(2, 2, 0)$-gain-tight gain graph which is connected and 4-regular. Suppose $v \in V_0$ has exactly four neighbours a, b, c, d. Then v is admissible if and only if the following conditions hold,*

(i) v is not contained in a balanced subgraph isomorphic to K_4, and,

(ii) v is not contained in a balanced subgraph isomorphic to $K_{1,1,3}$ with the property that v has degree 4 in this subgraph.

Proof. If (i) or (ii) fail to hold then v is clearly not admissible. For the converse, let (va, α), (vb, β), (vc, γ) and (vd, δ) be the edges incident to v.

We first show that if condition (ii) holds in Lemma 3.14, then c and d are not in H_{ab}. Suppose $N(v) \subset V(H_{ab})$. Since H_{ab} is balanced, we may assume, by switching, that all gains on H_{ab} are 1. Since all paths from a to b in H_{ab} have gain $\alpha\beta$, it follows that $\alpha = \beta$. If at least 3 edges incident to v all have gain 1 or all have gain -1, then we contradict $(2, 2)$-sparsity. Thus, without loss of generality, we may assume that $(va, 1), (vb, 1), (vc, -1), (vd, -1) \in E(G_0)$. Since G_0 is 4-regular, Lemma 3.3 implies that H_{ab} together with the edges incident to v is all of G_0. In particular H_{ab} was an induced subgraph of G_0. Thus it is easy to see that the reduction move at v adding $(ad, -1), (bc, -1)$ is admissible.

So without loss of generality, we may suppose that $d \notin V(H_{ab})$. We claim that $c \notin V(H_{ab})$. Suppose to the contrary that $c \in V(H_{ab})$. By the same argument as above, we may assume that $(va, 1), (vb, 1), (vc, -1) \in E(G_0)$. Notice that if we add the edges $(va, 1), (vb, 1), (vc, -1)$ to H_{ab}, then we obtain a subgraph H with $f(H) = 1$. Thus, H_{ab} is an induced subgraph of G_0, for otherwise

the fact that $d \notin V(H_{ab})$ would give a contradiction, by Lemma 3.3. Consider the reduction move at v adding $(bc, -1)$ and (ad, δ). If v is not admissible then one of the conditions in Lemma 3.14 must hold for the pair $(bc, -1)$ and (ad, δ). Since G_0 is 4-regular, Lemma 3.3 implies condition (i) of Lemma 3.14 cannot hold. If there is a balanced subgraph H_{ad} of $G_0 - v$ containing a, d with $f(H_{ad}) = 2$ in which all paths from a to d have gain δ, then, by Lemma 3.6, $H = H_{ab} \cup H_{ad}$ is balanced with $f(H) = 2$ and $N(v) \subset V(H)$. So we can repeat the argument used to show that $d \notin V(H_{ab})$ to obtain a contradiction. Suppose there is a balanced subgraph H_{bc} of $G_0 - v$ containing b, c with $f(H_{bc}) = 2$ in which all paths from b to c have gain -1. Then we must have $f(H_{ab} \cup H_{bc}) = f(H_{ab} \cap H_{bc}) = 2$, by Lemma 3.6. Hence $H_{ab} \cap H_{bc}$ is connected by Lemma 3.5 and every edge has gain 1, a contradiction. Similarly, condition (iii) of Lemma 3.14 cannot hold. Finally, the edge $(bc, -1)$ is clearly not in G_0, so we suppose that (ad, δ) is in G_0. By switching at the vertex d, we may assume that $\delta = 1$. We claim that the reduction move at v adding $(ac, -1), (bd, 1)$ is now admissible. If $(bd, 1)$ is present in G_0, then H_{ab} together with the 5 edges $(va, 1), (vb, 1), (vd, 1), (ad, 1), (bd, 1)$ violates $(2, 2)$-sparsity. So suppose there is a balanced subgraph H_{bd} of $G_0 - v$ containing b, d with $f(H_{bd}) = 2$ in which all paths from b to d have gain 1. But then, by Lemma 3.6, $H = H_{ab} \cup H_{bd}$ is balanced with $f(H) = 2$ and $N(v) \subset V(H)$, and hence the same argument from above may be applied once again to obtain a contradiction. If there is a balanced subgraph H_{ac} of $G_0 - v$ containing a, c with $f(H_{ac}) = 2$ in which all paths from a to c have gain -1, then $f(H_{ab} \cup H_{ac}) = f(H_{ab} \cap H_{ac}) = 2$, by Lemma 3.6. Hence $H_{ab} \cap H_{ac}$ is connected and every edge has gain 1, a contradiction. Similarly, condition (iii) of Lemma 3.14 cannot hold for the pair $(ac, -1), (bd, 1)$. So the reduction move at v adding $(ac, -1), (bd, 1)$ is indeed admissible. Thus, we conclude that $c \notin V(H_{ab})$.

Now, suppose that conditions (i) and (ii) in the statement of this lemma hold. Then condition (iv) in Lemma 3.14 cannot hold for all possible pairs of edges, so we may assume that the edges $(ab, \alpha\beta), (cd, \gamma\delta)$ do not exist in (G_0, ψ). Since G_0 is 4-regular, Lemma 3.3 implies that condition (i) in Lemma 3.14 does not hold. Suppose H is a balanced subgraph of (G_0, ψ) that satisfies condition (iii) in Lemma 3.14. By switching, we may assume that all gains on H are 1. By the assumptions in condition (iii) it follows that $\alpha\beta = \gamma\delta = 1$, contradicting $(2, 2)$-sparsity of G_0 if $\alpha = \beta = \gamma = \delta$. So suppose without loss of generality that $\alpha = \beta = 1$ and $\gamma = \delta = -1$. Note that at most one of the edges $(ac, -1), (ad, -1), (bc, -1), (bd, -1)$ can be present in G_0, for otherwise the graph obtained from H by adding two of those edges together with the four edges incident to v violates $(2, 0)$-sparsity. Hence $(ad, -1), (bc, -1) \notin E(G_0)$ or $(ac, -1), (bd, -1) \notin E(G_0)$. Without loss of generality we assume that $(ad, -1), (bc, -1) \notin E(G_0)$ and consider the reduction operation on v that adds $(ad, -1)$ and $(bc, -1)$.

Suppose there exists a balanced subgraph H' of (G_0, ψ) that satisfies condition (iii) in Lemma 3.14 for the pair $(ad, -1), (bc, -1)$, that is, all paths from a to d and all paths from b to c in H' have gain -1. Then $H \cap H'$ is disconnected (since there cannot be any path from a to d or from b to c in $H \cap H'$). It follows that $f(H \cap H') \geq 2 + 2 = 4$ and hence $f(H \cup H') \leq 3 + 3 - 4 = 2$. Thus, $f(H \cup H') = 2$, for otherwise the graph obtained from $H \cup H'$ by adding the four edges incident to v would violate $(2, 0)$-sparsity. This says that $H \cap H'$ has exactly two components C_1 and C_2, with either $a, b \in V(C_1)$ and $c, d \in V(C_2)$ or $a, c \in V(C_1)$ and $b, d \in V(C_2)$.

Suppose first that $a, b \in V(C_1)$ and $c, d \in V(C_2)$. Then every path in H' from a to b and from c to d has gain 1. Since every path from a to d and from b to c has gain -1 in H', and H' is balanced,

it follows that every path from a to c and from b to d must have gain -1. Therefore, H' together with the four edges incident to v is balanced and violates $(2,2)$-sparsity.

Suppose next that $a, c \in V(C_1)$ and $b, d \in V(C_2)$. Then every path in H' from a to c and from b to d has gain 1. This implies that $(ac, -1)$ and $(bd, -1)$ cannot exist in G_0. (If one of those edges did exist in G_0, then it cannot be an edge of $H \cup H'$, and hence $H \cup H'$ together with this edge and the four edges incident to v would violate $(2,0)$-sparsity.) We now consider the reduction operation on v that adds $(ac, -1)$ and $(bd, -1)$. Suppose there exists a balanced subgraph H'' of (G_0, ψ) that satisfies condition (iii) in Lemma 3.14 for the pair $(ac, -1)$, $(bd, -1)$, that is, all paths from a to c and all paths from b to d in H'' have gain -1. Note that $(H \cup H') \cap H''$ is disconnected since there cannot be any path from a to c or from b to d in $(H \cup H') \cap H''$. Therefore, $f((H \cup H') \cap H'') \geq 2 + 2 = 4$ and hence $f((H \cup H') \cup H'') \leq 2 + 3 - 4 = 1$. So the subgraph obtained from $(H \cup H') \cup H''$ by adding the four edges incident to v violates $(2,0)$-sparsity. So if the reduction operation on v that adds $(ac, -1)$ and $(bd, -1)$ is not admissible, then condition (ii) in Lemma 3.14 holds for $(ac, -1)$, $(bd, -1)$. As we have shown in the beginning of this proof, we may assume that there exists a balanced subgraph H_{ac} of (G_0, ψ) containing a, c but not b, d, v with $f(H_{ac}) = 2$ and every path from a to c in H_{ac} has gain -1. Then $H \cap H_{ac}$ is disconnected (since there cannot be a path from a to c in $H \cap H_{ac}$) and hence $f(H \cup H_{ac}) \leq 3 + 2 - 4 = 1$. It follows that $H \cup H_{ac}$ together with the four edges incident to v violates the $(2,0)$-sparsity of G_0.

Therefore, if v is not admissible, then condition (ii) in Lemma 3.14 holds for $(ad, -1)$, $(bc, -1)$. As we have shown in the beginning of this proof, we may assume that there exists a balanced subgraph H_{bc} of (G_0, ψ) containing b, c but not a, d, v with $f(H_{bc}) = 2$ and every path from b to c in H_{bc} has gain -1. But then, by the same argument as in the paragraph above, we have $f(H \cup H_{bc}) \leq 3 + 2 - 4 = 1$, which contradicts the $(2,0)$-sparsity of G_0.

So if v is not admissible, then condition (ii) in Lemma 3.14 holds for $(ab, \alpha\beta), (cd, \gamma\delta)$. By using the argument in the beginning of this proof again, we may assume that there exists a balanced subgraph H_{ab} of (G_0, ψ) containing a, b but not c, d, v with $f(H_{ab}) = 2$ and every path from a to b in H_{ab} has gain $\alpha\beta$. Consider the reduction operation on v that adds $(ad, \alpha\delta)$ and $(bc, \beta\gamma)$.

If this move is not admissible then one of the conditions (ii), (iii) or (iv) in Lemma 3.14 holds for $(ad, \alpha\delta), (bc, \beta\gamma)$. Suppose first that (iv) fails, that is, $(ad, \alpha\delta)$ and $(bc, \beta\gamma)$ are not edges of (G_0, ψ). Suppose there exists a balanced subgraph H of (G_0, ψ) that satisfies condition (iii) for the pair $(ad, \alpha\delta), (bc, \beta\gamma)$, that is, all paths from a to d in H have gain $\alpha\delta$ and all paths from b to c in H have gain $\beta\gamma$. If $H \cap H_{ab}$ is connected, then every path from a to b in H has gain $\alpha\beta$. But since H is balanced, this implies that the subgraph H' of G_0 consisting of H and the four edges incident to v is balanced and satisfies $f(H') = 1$, contradicting $(2,2)$-sparsity. Thus $H \cap H_{ab}$ is disconnected, and hence $f(H \cup H_{ab}) \leq 3 + 2 - 4 = 1$. But since $H \cup H_{ab}$ contains all four neighbours of v, this contradicts $(2,0)$-sparsity. So we may assume that condition (ii) in Lemma 3.14 holds for $(ab, \alpha\beta), (cd, \gamma\delta)$. Without loss of generality we may assume that there exists a balanced subgraph H_{ad} of (G_0, ψ) containing a, d but not b, c, v with $f(H_{ad}) = 2$ and every path from a to d in H_{ad} has gain $\alpha\delta$. If $H_{ad} \cap H_{ab}$ is connected, then $H_{ad} \cup H_{ab}$ is balanced with $f(H_{ad} \cup H_{ab}) = 2$, and the subgraph H' of (G_0, ψ) consisting of $H_{ad} \cup H_{ab}$ and the three edges joining v with a, b and d is also balanced with $f(H') = 1$, contradicting $(2,2)$-sparsity. Thus $H_{ad} \cap H_{ab}$ is disconnected, and hence $f(H_{ad} \cup H_{ab}) \leq 2 + 2 - 4 = 0$. But since $H_{ad} \cup H_{ab}$ contains three neighbours of v, this contradicts $(2,0)$-sparsity.

So without loss of generality we may assume that $(bc, \beta\gamma)$ is an edge of (G_0, ψ). Then the edge $(bd, \beta\delta)$ cannot exist in G_0 (for otherwise b would have degree at most 1 in H_{ab} by 4-regularity of G_0, a contradiction). The edge $(ac, \alpha\gamma)$ can also not exist in G_0, for otherwise the graph consisting of H_{ab} and the edges $(va, \alpha), (vb, \beta), (vc, \gamma), (ac, \alpha\gamma), (bc, \beta\gamma)$ would be balanced and would violate $(2, 2)$-sparsity. So consider the reduction operation on v that adds $(ac, \alpha\gamma)$ and $(bd, \beta\delta)$. If there exists a balanced subgraph H of (G_0, ψ) that satisfies condition (iii) for the pair $(ac, \alpha\gamma), (bd, \beta\delta)$, then, by the same argument as in the paragraph above, $H \cap H_{ab}$ is disconnected, and hence $f(H \cup H_{ab}) \leq 3 + 2 - 4 = 1$. But since $H \cup H_{ab}$ contains all four neighbours of v, this contradicts $(2, 0)$-sparsity. So we may assume that condition (ii) in Lemma 3.14 holds for $(ac, \alpha\gamma), (bd, \beta\delta)$. Without loss of generality we may assume that there exists a balanced subgraph H_{ac} of (G_0, ψ) containing a, c but not b, d, v with $f(H_{ac}) = 2$ and every path from a to c in H_{ac} has gain $\alpha\gamma$. As in the paragraph above, we see that $H_{ac} \cap H_{ab}$ is disconnected, and hence $f(H_{ac} \cup H_{ab}) \leq 2+2-4 = 0$. But since $H_{ac} \cup H_{ab}$ contains three neighbours of v, this contradicts $(2, 0)$-sparsity. □

3.5. Graph contractions

We now consider the existence of suitable triangles or K_4's in order to apply the reverse vertex splitting move or the reverse vertex-to-K_4 move. After giving conditions on when they can be applied we will also prove a couple of technical lemmas needed in the next section.

Lemma 3.16. *Let (G_0, ψ) be a $(2, 2, 0)$-gain-tight gain graph. Suppose (G_0, ψ) contains a balanced subgraph K isomorphic to K_4 which induces at most one additional edge. Then a reverse vertex-to-K_4 move at K is admissible unless there is a vertex x and edges $(xa, \alpha), (xb, \alpha)$ for some $a, b \in V(K)$.*

Proof. Let K' denote the contraction of the graph K^* induced by K. (So K' is either a single vertex or a single vertex with a loop). Then $f(K^*) = f(K')$ and hence the lemma follows from a simple counting argument. □

Lemma 3.17. *Let (G_0, ψ) be a $(2, 2, 0)$-gain-tight gain graph. Suppose (G_0, ψ) contains a subgraph K isomorphic to K_3 with $V(K) = \{a, b, c\}$ and with all three edges of K having gain 1. Suppose the edge $(ab, -1)$ does not exist in G_0. Then a reverse vertex-splitting at K contracting the edge $(ab, 1)$ is non-admissible if and only if one of the following conditions holds:*

(i) there is a subgraph H_0 of G_0 containing a, b and the edge $(ab, 1)$, but not c, with $f(H_0) = 0$;

(ii) there is a balanced copy of K_3 containing a, b and some vertex $d \neq c$;

(iii) there is a balanced subgraph H_0 of G_0 containing a, b and the edge $(ab, 1)$ with $f(H_0) = 2$, and if H_0 contains c then it does not contain the edges $(ca, 1)$ and $(cb, 1)$;

(iv) there are loops incident to both a and b, or both edges $(ac, -1)$ and $(bc, -1)$ exist.

Proof. If any one of these conditions holds then it is clear that v is not admissible. Conversely, suppose the reverse vertex-splitting at K contracting the edge $(ab, 1)$ is non-admissible. Let (G_0', ψ') be the gain graph resulting from this reverse vertex-splitting move and let $\overline{a}$ be the vertex of (G_0', ψ') corresponding to the vertex pair a and b in (G_0, ψ). Then either the covering graph of (G_0', ψ')

is not simple, or there is a subgraph H_0' of (G_0', ψ') with $f(H_0') < 0$, or there is a balanced subgraph of (G_0', ψ') with $f(H_0') < 2$. In the first case (ii) or (iv) holds. In the second case, H_0' clearly contains $\overline{a}$, but it cannot contain c, for otherwise $f(H_0) = f(H_0') < 0$ if $(c\overline{a}, 1) \in E(H_0')$ or $f(H_0 \cup \{(ca, 1)\}) = f(H_0') < 0$ if $(c\overline{a}, 1) \notin E(H_0')$, where H_0 is the subgraph of (G_0, ψ) that is obtained from H_0' by the vertex splitting move at $\overline{a}$. Thus, H_0 contains a, b and the edge $(ab, 1)$, but not c, and satisfies $f(H_0) = 0$. In the third case, the balanced H_0' again clearly contains $\overline{a}$. If it contains c and $(c\overline{a}, 1)$ then $f(H_0) = f(H_0') < 2$, a contradiction. So if H_0' contains c then it does not contain $(c\overline{a}, 1)$. Thus, H_0 contains a, b and the edge $(ab, 1)$, and if H_0 contains c then it does not contain the edges ca and cb. Moreover, H_0 is balanced with $f(H_0) = 2$. □

We now follow the approach in [12]. Define a *triangle sequence* $T_1, T_2, \ldots, T_n$ where T_1 is a triangle on vertices a, b, c and T_{i+1} is formed from T_i by adding a vertex of degree 2 adjacent to two vertices x, y of T_i such that $xy \in E(T_i)$ and x, y are in exactly one triangle in T_i. A triangle sequence is *balanced* if each T_i (or equivalently just T_n) is balanced. A *maximal balanced triangle sequence* is a balanced triangle sequence that cannot be extended to a larger balanced triangle sequence. A *chord* of T_n is an edge in the subgraph $G_0[V(T_n)]$ of G_0 induced by $V(T_n)$ which is not in T_n.

The following lemma is easy to deduce from the definitions.

Lemma 3.18. *Let (G_0, ψ) be $(2, 2, 0)$-gain-tight gain graph and let T_n be a balanced triangle sequence whose edges all have gain 1. Then $f(T_n) = 3$ and T_n has at most 3 chords. Moreover, if three chords exist, at least two of them have gain -1.*

In the next lemma we show that there exists an edge for which conditions (ii) and (iv) of Lemma 3.17 do not hold, provided that the maximal balanced triangle sequence has enough vertices.

Lemma 3.19. *Let (G_0, ψ) be $(2, 2, 0)$-gain-tight and let T_n be a maximal balanced triangle sequence. Suppose $|V(T_n)| \geq 6$. Then there exists an edge of T_n with the properties that it does not have a parallel edge, its end vertices are not both incident to a loop, and it is contained in exactly one balanced triangle in (G_0, ψ).*

Proof. Since T_n is balanced we may assume that the gain on every edge of T_n is 1. Let $s_1, s_2, \ldots, s_r$ be the edges of T_n contained in exactly one triangle in T_n.

We first show that the s_i form a simple cycle spanning $V(T_n)$. This can be verified by induction. It clearly holds for $n = 1$. Suppose it holds for all $m < n$ and consider T_n. Let T_n be formed from T_{n-1} by adding a triangle on a, b, c such that $a, b \in V(T_{n-1})$ and $c \notin V(T_{n-1})$. Then by induction there is a simple spanning cycle C in T_{n-1} with $ab \in C$. We construct the simple spanning cycle for T_n by removing the edge ab from C and adding the edges ac, cb to C.

By Lemma 3.18, T_n has at most 3 chords, and at most one them has gain 1. Let C be the simple cycle spanning $V(T_n)$ consisting of the edges $s_1, s_2, \ldots, s_r$. Since $V(C)$ has at least 6 vertices, C has at least 6 edges. We say that an edge of C is *blocked* if it has a parallel edge, or its endvertices are both incident to a loop, or it is contained in more than one balanced triangle in (G_0, ψ). We claim that C contains at least one edge that is not blocked.

Note that a chord with gain 1 can block at most three edges of C. A chord with gain -1 can block at most one edge of C, since it cannot create a balanced triangle on its own, but it could be parallel to an edge of C. Further, two non-loop chords with gain -1 only create a balanced triangle

if they are of the form $(cd, -1)$ and $(ce, -1)$ for some $c, d, e \in V(T_n)$. In this case at most one edge of C is contained in a triangle with these two chords. Thus, two non-loop chords with gain -1 can block at most two edges of C. Similarly, three non-loop chords with gain -1 can block at most three edges of C. Finally, two chords that are loops can block at most one edge of C, and three chords that are loops can block at most two edges of C.

So if at least two of the three chords are loops, then there are at most 4 blocked edges in C. So suppose exactly one of the chords is a loop. Since at most one of the remaining two chords can have gain 1, it follows that there are again at most 4 blocked edges in C. Finally, if none of the chords is a loop, then there are at most 5 blocked edges in C. Thus there exists an edge of C that is not blocked, as claimed. □

3.6. The inductive construction

We can now put together our results to prove the desired characterisation of $(2, 2, 0)$-gain-tight gain graphs.

Theorem 3.20. *Let (G_0, ψ) be a gain graph. Then (G_0, ψ) is $(2, 2, 0)$-gain-tight if and only if (G_0, ψ) can be generated from vertex disjoint copies of graphs in $\mathcal{B}$ by applying H1, H2, H3, vertex-to-K_4 and vertex splitting moves.*

Proof. The easy direction is given by Lemma 3.7. For the converse, observe first that if (G_0, ψ) is disconnected then every connected component of (G_0, ψ) is $(2, 2, 0)$-gain-tight. So we will prove that an arbitrary *connected* $(2, 2, 0)$-gain-tight gain graph has an admissible reverse move which results in another (possibly disconnected) $(2, 2, 0)$-gain-tight gain graph with fewer vertices, or is one of the base graphs in $\mathcal{B}$. The theorem then follows by induction on $|V_0|$.

So from now on we will assume that (G_0, ψ) is connected. Observe that (G_0, ψ) is either 4-regular or contains a vertex of degree 2 or 3. We consider the following cases.

Case 1. G_0 contains a vertex with two incident edges or a degree 3 vertex with exactly two neighbours, which is not contained in a subgraph isomorphic to R.

Lemmas 3.8 and 3.10 show that there exists an admissible reverse move.

We can now assume that Case 1 does not hold. In the following, we let $\hat{K}_4$ be the $(2, 2, 0)$-gain-tight gain graph consisting of a balanced K_4 as well as one additional vertex x and the four edges $(xa, 1), (xa, -1), (xb, 1), (xb, -1)$ where a and b are distinct vertices of the K_4.

Case 2.a. G_0 contains a degree 3 vertex v with exactly 3 neighbours and v is not contained in a subgraph isomorphic to K_4^+ or $\hat{K}_4$.

Lemma 3.9 implies v is admissible or v is contained in a balanced subgraph K isomorphic to K_4. By Lemma 2.17 we may assume every edge of K has gain 1. Note that K is an induced subgraph of G_0. If K is not admissible to contract then Lemma 3.16 implies there is a vertex x and edges $(xa, \alpha), (xb, \alpha)$ for $a, b \in V(K)$ and $\alpha \in \{1, -1\}$. We may apply Lemma 2.17 again at x

to make $\alpha = 1$. Now we have a balanced K_3 on x, a, b. We will denote a balanced K_3 on vertices r, s, t by $K_3(r, s, t)$.

Consider a maximal balanced triangle sequence $T_1 = K_3(v, a, c), T_2 = T_1 \cup K_3(a, b, c), T_3 = T_2 \cup K_3(a, b, x), \ldots, T_n$ (where c is the final vertex of K). By Lemma 2.17 we may assume every edge of T_n has gain 1. By Lemma 3.17 an edge $(rs, 1)$ in a $K_3(r, s, t)$ in T_n, with $(rs, -1)$ not an edge of G_0, is non-admissible in G_0 for a reverse vertex-splitting move if and only if condition (i), (ii), (iii) or (iv) holds.

We claim that there exists an edge in T_n for which (ii) and (iv) does not hold. This follows from Lemma 3.19 if $|V(T_n)| \geq 6$. If $|V(T_n)| = 5$, then T_n is uniquely determined, and since K is an induced subgraph of G_0, and v is not contained in a subgraph isomorphic to $\hat{K}_4$, the existence of such an edge in T_n can easily be verified by inspection.

So let $(rs, 1)$ be an edge of T_n for which (ii) and (iv) do not hold. Let t be the third vertex of the unique balanced triangle containing $(rs, 1)$. If there is a subgraph H_0 satisfying (i) then we contradict the $(2, 0)$-sparsity of G_0 as follows. Since t has two neighbours (namely r and s) in H_0, clearly any other neighbour of t in T_n is not in H_0. Note that T_n contains a balanced triangle of the form $K_3(r, t, w)$ or of the form $K_3(s, t, w')$ (or both). Suppose that $K_3(r, t, w)$ exists in T_n. Then we may repeat the above argument for w and the subgraph H_0' which is obtained from H_0 by adding t and the edges $(tr, 1)$ and $(ts, 1)$ to see that any neighbour of w in T_n (other than r and t) is not in H_0'. By iterating this argument we conclude that none of the vertices of $V(T_n) \setminus \{r, s\}$ are in H_0. It now follows that $H_0 \cup T_n \cup (vb, 1)$ violates $(2, 0)$-sparsity. Thus, (i) also does not hold for $(rs, 1)$.

Finally, we show that there is no balanced subgraph H_0 satisfying (iii) for the edge $(rs, 1)$. Suppose to the contrary that there does exist such a subgraph H_0.

If $t \notin V(H_0)$ then we claim that none of the neighbours of t in T_n (except for r and s) are in H_0. To see this, we suppose for a contradiction (with the notation of the previous paragraph) that $w \in V(H_0)$. Then either $(wr, 1) \in E(H_0)$ or $(wr, 1) \notin E(H_0)$. In the first case adding t and its 3 incident edges to H_0 contradicts $(2, 2, 0)$-gain-sparsity, so the latter case must hold. In this case, there exists a path from r to w in H_0 which has gain -1, for otherwise adding the edge $(wr, 1)$ to H_0 yields a graph that violates $(2, 2, 0)$-gain-sparsity. By the balancedness of H_0, all paths from r to w in H_0 have gain -1. Since all edges of T_n have gain 1, it follows that $T_n \cap H_0$ is disconnected, with r and w being in different connected components. If $T_n \cap H_0$ has more than two connected components, or if any edge incident to w is in $T_n \cap H_0$, then, since $(rs, 1) \in T_n \cap H_0$, we have $f(T_n \cap H_0) \geq 3 + 3 = 6$. In this case we have $f(T_n \cup H_0) \leq 3 + 2 - 6 < 0$, a contradiction. So suppose $T_n \cap H_0$ has exactly two connected components and one of them consists only of w. Then we only have $f(T_n \cap H_0) \geq 3 + 2 = 5$ and $f(T_n \cup H_0) \leq 3 + 2 - 5 = 0$. We claim that in this case the edge $(vb, 1)$ is not in H_0. Suppose to the contrary that $(vb, 1) \in E(H_0)$. Then $(rs, 1)$, v and b must belong to the same connected component of $T_n \cap H_0$, and hence there either exists a vertex $y \in T_n \cap H_0$, $y \neq s, w$, such that $(ry, 1) \in T_n \cap H_0$, or there exists a vertex $z \in T_n \cap H_0$, $z \neq r, w$, such that $(sz, 1) \in T_n \cap H_0$. In the first case, we must have $(wy, 1) \in T_n \setminus H_0$, and in the second case we must have $(tz, 1) \in T_n \setminus H_0$. Thus, by adding the vertex t and the edges $(tr, 1), (ts, 1), (wr, 1), (wt, 1)$ together with the fifth edge $(wy, 1)$ ($(tz, 1)$, respectively) to H_0 we obtain a graph that violates $(2, 2, 0)$-gain-sparsity. Hence $(vb, 1)$ is not in H_0 and $T_n \cup H_0 \cup (vb, 1)$ violates $(2, 2, 0)$-gain-sparsity. Thus, as claimed, any neighbour of t in T_n (other than r and s) is not in H_0. We may now continue this argument as in the paragraph above to show that none of the vertices of $V(T_n) \setminus \{r, s\}$

are in H_0. It follows that $T_n \cup H_0 \cup (vb, 1)$ violates $(2, 2, 0)$-gain-sparsity.

So we may assume that $t \in V(H_0)$. Then, by condition (iii), the edges $(tr, 1), (ts, 1)$ are not in H_0. The maximal balanced triangle sequence T_n contains a balanced triangle of the form $K_3(r, t, w)$ or of the form $K_3(s, t, w')$ (or both). As above, we assume that $K_3(r, t, w)$ exists. Suppose first that $w \in V(H_0)$. Then the edges $(wt, 1)$ and $(wr, 1)$ are also both in H_0. (If not, say $(wt, 1) \notin E(H_0)$, then the graph obtained from H_0 by adding the edges $(tr, 1), (ts, 1), (wt, 1)$ violates $(2, 0)$-sparsity.) Since the edges $(wt, 1)$ and $(wr, 1)$ form a path of gain 1 from r to t in H_0, and H_0 is balanced, every path from r to t in H_0 has gain 1. Thus, if we add the edge $(tr, 1)$ to H_0, then we obtain a balanced graph that violates the $(2, 2)$-sparsity of G_0, a contradiction.

So we may assume that $w \notin V(H_0)$. Let H_0' be the graph obtained from H_0 by adding the edges $(tr, 1)$ and $(ts, 1)$. Then $f(H_0') = 0$. Since w has two neighbours (namely r and t) in H_0', clearly any other neighbour of w in T_n is not in H_0'. Using again the same iteration argument from above, we conclude that none of the vertices of $V(T_n) \setminus \{r, s, t\}$ are in H_0'. It now follows that $H_0' \cup T_n \cup (vb, 1)$ violates $(2, 0)$-sparsity.

Thus, in Case 2a there exists an admissible reverse move.

Case 2.b. G_0 contains a degree 3 vertex, and every degree 3 vertex in G_0 is contained in an induced subgraph isomorphic to R, K_4^+, K_4^{++} or $\hat{K}_4$.

Every vertex of degree 3 is contained in an induced subgraph W_i isomorphic to K_4^+ or in an induced subgraph Z_j isomorphic to either R, $\hat{K}_4$, or a graph K_4^{++} from Figures 1 (d)-(g). Since $f(K_4^+) = 1$ and $f(R) = f(\hat{K}_4) = f(K_4^{++}) = 0$ we say the former are type 1 and the latter are type 0. Let $W_1, \ldots, W_{r'}$ be all such type 1 induced subgraphs and $Z_1, \ldots, Z_{s'}$ be all such type 0 induced subgraphs.

Note that for all $1 \leq i \leq s'$, any proper non-empty subgraph H of Z_i has $f(H) \geq 1$. Now, for any pair of subgraphs Z_i, Z_j we have that Z_i and Z_j are necessarily vertex disjoint. If not then $f(Z_i \cap Z_j) > 0$ and hence $f(Z_i \cup Z_j) < 0$, which would contradict the $(2, 2, 0)$-sparsity of G_0. Next, for all $1 \leq j \leq r'$, any proper non-empty subgraph Y of W_j is either a loop or has $f(Y) \geq 2$. Thus, for any pair of subgraphs Z_i, W_j we have that either the intersection is a loop, or Z_i and W_j are vertex disjoint. If not then $f(Z_i \cap W_j) \geq 2$ and hence $f(Z_i \cup W_j) = 1 - f(Z_i \cap W_j) < 0$, contradicting $(2, 2, 0)$-sparsity. Lastly, for any pair of subgraphs W_i, W_j with non-empty intersection, we must have $f(W_i \cap W_j) \in \{1, 2\}$. This implies that $W_i \cap W_j$ is either empty, a loop, or has $f(W_i \cap W_j) = 2$. Moreover in the case when $f(W_i \cap W_j) = 2$, then $W_i \cap W_j$ is either a double edge or a single vertex. (The case when $W_i \cap W_j$ is a copy of K_4 would mean either $i = j$ or would contradict the fact that W_i and W_j are induced subgraphs of G_0. Similarly, the cases when $W_i \cap W_j$ is a triangle plus an edge, or an edge plus a loop would contradict the $(2, 2, 0)$-gain sparsity of G_0.)

We next modify our lists $W_1, \ldots, W_r$ and $Z_1, \ldots, Z_s$ until any pair W_i, W_j and any pair Z_i, W_j are vertex disjoint. First, whenever a pair of W_i and W_j intersect in a double edge or a single vertex, then we discard them and add the union of W_i and W_j as a new Z_m. If this process is iterated, then the new list of W_i's and Z_i's has the property that any two elements in the list either do not intersect or intersect in a loop. Whenever a pair of W_i and W_j in the list intersect in a loop, then we discard them and add the union of W_i and W_j as a new W_ℓ, and whenever a pair of W_i and Z_j intersect in a loop then we discard them and add the union of W_i and Z_j as a new Z_k. This process is iterated

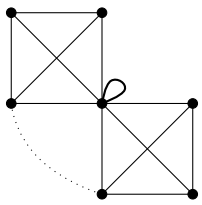

Figure 6: The graph consisting of some $k \geq 2$ copies of K_4 which all intersect in a single vertex and that vertex is incident with one loop. This family of graphs gives the only additional isomorphism classes that can occur in the modified list of W_i's. Gain labels omitted.

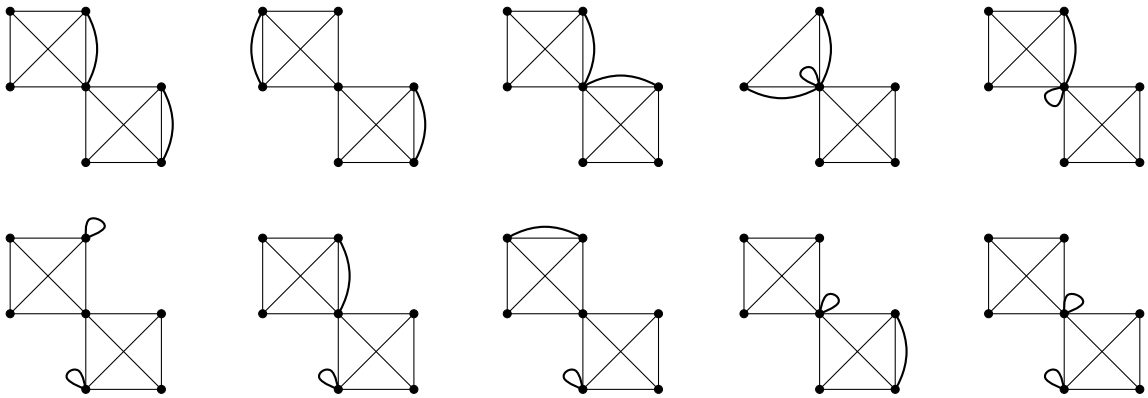

Figure 7: Additional isomorphism classes that can occur in the first step of the modified list of Z_j's. Gain labels omitted.

until the final list of all W_i's and Z_i's has the desired property that any two elements in the list are vertex disjoint. (The additional W_i's at any step of the process and the additional Z_i's from step 1 of this process are as depicted in Figures 6 and Figures 7 and 8.)

Let U and F be the sets of vertices and edges of G_0 which are in none of the W_i and in none of the Z_j. Associate with G_0 an auxiliary (multi)graph G_0^* which has a vertex for each W_i, a vertex for each Z_i and a vertex for U and has an edge corresponding to each edge of G_0 of the form $x_i x_j$, where x_i, x_j are taken from distinct elements of $V(G_0^*) = \{W_1, \ldots, W_r, Z_1, \ldots, Z_s, U\}$. Also define G_0^- to be the simple graph which is obtained from G_0^* by removing any parallel edges.

The connectivity of G_0 implies that G_0^- is connected. Suppose $|V(G_0^-)| = 1$. Then G_0 is a copy of some Z_j. If $G_0 = \hat{K}_4$ then we may apply a reverse H3d move. If G_0 is a union of two K_4^+ graphs which intersect in a double edge then we may apply a reverse vertex-to-K_4 move (see Figure 8). If G_0 is not a base graph then by the iteration process above, it follows in all remaining cases that G_0 contains a cut-vertex that separates an induced subgraph isomorphic to K_4^+ from the rest of the graph. It is now clear that we may use Lemma 3.16 to apply a reverse vertex-to-K_4 move to a K_4^+ subgraph of G_0. Thus, we may suppose that $|V(G_0^-)| > 1$.

Suppose $r = 0$. Let $G_0[U]$ denote the subgraph of G_0 induced by U. Note that since $f(Z_i) = 0$ for each i, no two Z_i can be adjacent (by the $(2, 2, 0)$-gain-sparsity of G_0). Similarly, there is at most one edge between a W_i and a Z_j and at most two edges between a W_i and a W_k. Since G_0^- is

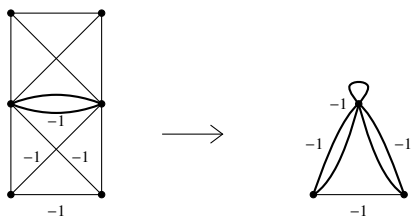

Figure 8: An additional isomorphism class that can occur in the first step of the modified list of Z_j's and the result, R, of a reverse vertex-to-K_4 move applied to this graph. Omitted gain labels are equal to 1.

connected, it follows that G_0^- is the graph $K_{1,s}$ where $s \geq 1$. Moreover, we have

$$f(G_0[U]) - d(U, V_0 - U) = \sum_{i=1}^{s} f(Z_i) + f(G_0[U]) - d(U, V_0 - U) = f(G_0) = 0.$$

Since every vertex in U has degree at least 4 in G_0, by Lemma 3.4 we have,

$$d(U, V_0 - U) \geq 2f(G_0[U]) = 2d(U, V_0 - U).$$

Thus $d(U, V_0 - U) = 0$, a contradiction.

Now suppose $r > 0$. Recall that each W_i is a K_4^+ or is of the form illustrated in Figure 6. Hence, if any W_i is not incident to two parallel edges in G_0^* then we may contract a copy of K_4^+ to a loop by Lemma 3.16. So we suppose that every W_i is incident to two parallel edges in G_0^*.

We calculate

$$f(G_0) = \sum_{i=1}^{r} f(W_i) + \sum_{j=1}^{s} f(Z_j) + 2|U| - |F|, \tag{4}$$

which implies that $|F| = 2|U| + r$.

Suppose first that every W_i and every Z_j is incident to at least two edges in F. Since each vertex in U has degree at least 4, there are at least $4|U| + 2(r + s)$ edge/vertex incidences in F. This implies $|F| \geq 2|U| + r + s$, and hence $s = 0$. By a similar counting argument, if some W_i is incident to more than two edges in F then there are at least $4|U| + 2(r - 1) + 3$ edge/vertex incidences in F. This implies $|F| > 2|U| + r$, a contradiction. Thus each W_i has degree exactly 2 in G_0^*. This implies that either G_0 is the disjoint union of W_1 and W_2 with two edges between them, or, G_0^- is the graph $K_{1,r}$. In the former case, there is an admissible reverse vertex-to-K_4 move which contracts a copy of K_4^+ to a loop. So suppose G_0^- is $K_{1,r}$. In this case, every $v \in U$ has degree exactly 4. We may assume that every W_i is a copy of K_4^+ and that for every W_i there exists a vertex in U that is joined to two vertices in W_i by edges with identical gains. (Otherwise, there is an admissible reverse vertex-to-K_4 move.)

Let $v \in U$ be adjacent to two vertices in some W_i. Then there is balanced copy of K_3 containing v and two vertices a, b of W_i. We may assume the gains on this copy of K_3 are all 1. We now apply Lemma 3.17 to show that there is an admissible reverse vertex-splitting move that contracts either $(va, 1)$ or $(vb, 1)$. Since W_i is a copy of K_4^+ at most one of a, b can have a loop. We suppose there is no loop on a and consider the contraction of $(va, 1)$. It follows from Lemma 3.17 that if $(va, 1)$

is non-admissible then there is either a subgraph H_0 containing v, a and the edge $(va, 1)$ but not b with $f(H_0) = 0$ or a balanced subgraph H_0 containing v, a and the edge $(va, 1)$ with $f(H_0) = 2$ and if H_0 contains b then it does not contain the edges $(ba, 1)$ and $(vb, 1)$. In both cases it is easy to deduce from the structure of G_0 that such an H_0 cannot exist.

So we suppose that there is a Z_j which is incident to only one edge in F. If a W_i is joined to another W_k by two edges, then we discard them and add the union of W_i and W_k together with the two extra edges as a new Z_m. Similarly, if a W_i is joined to a Z_k by one edge, then we discard them and add the union of W_i and Z_k together with the extra edge as a new Z_ℓ. Note that these new Z_i cannot be adjacent for the same reasons as before. Then the graph G_0^- corresponding to this new underlying structure of G_0 contains the graph $K_{1,t}$ as a spanning subgraph, where the vertices in the partite set of size $t = b + c$ correspond to the graphs $W_1, \ldots, W_b$ and $Z_1, \ldots, Z_c$, and the vertex in the other partite set corresponds to U. If G_0^- is not equal to $K_{1,t}$, then the additional edges must join vertices corresponding to the W_i, and these edges represent single edges among pairs of W_i in G_0^*. We call this set of edges A. Note that in G_0^*, for each W_i there are two parallel edges joining W_i to U, and for each Z_j there is one edge joining Z_j to U. There may be additional edges joining a W_i with U or a Z_j with U, and we denote this set of additional edges by B. Hence we have $f(G_0[U]) = b+c+d$, where $d = |A|+|B|$. Thus, by Lemma 3.4, we have $d(U, V_0 - U) \geq 2(b+c+d)$. But there are only exactly $2b + c + d - |A|$ edges incident to U, so we have $|A| = c = d = 0$, and $c = 0$ contradicts our assumption that there exists a Z_j.

Thus, in Case 2b there exists an admissible reverse move.

Henceforth we may assume that G_0 is 4-regular. First let us deal with two special possible subgraphs of G_0.

Case 3.a. There exists a balanced subgraph isomorphic to either a balanced copy of $K_{1,1,3}$, or, to a balanced copy of K_4 which neither induces additional edges nor is contained in a copy of $K_5 - e$.

Suppose that G_0 contains an induced balanced copy of K_4. By switching we may assume that all edges of the K_4 have gain 1. By Lemma 3.16 and the assumptions of this case, if a reverse vertex-to-K_4 move is not admissible then there is a vertex s adjacent to exactly two of the K_4 vertices such that $K_4 \cup s$ is balanced. Consider the vertex s in $K_4 \cup s$. Since G_0 is 4-regular and K_4 is not contained in $K_5 - e$, s has either a neighbour not in the K_4 with two parallel edges to s, or two neighbours not in the K_4, or a loop on s. (These three possibilities are illustrated in Figure 9(a).) In the first case we can use Lemma 3.13 to conclude that s is admissible and in the second case we can use Lemma 3.15 to conclude that s is admissible. In the third case we may use Lemma 3.17. Let r, t be the neighbours of s in G_0. By switching, we may assume that $(sr, 1)$ and $(st, 1)$ are edges in G_0. Consider the contraction of the edge $(sr, 1)$. Since G_0 is 4-regular, conditions (ii) and (iv) of Lemma 3.17 evidently fail. Condition (i) fails by Lemma 3.3. Finally, since G_0 is 4-regular and any balanced subgraph H_0 with $f(H_0) = 2$ cannot have a vertex of degree 1, condition (iii) also fails.

Suppose then that G_0 contains a balanced copy of $K_{1,1,3}$ and does not contain an induced balanced copy of K_4. By switching, we may assume that all edges of $K_{1,1,3}$ have gain 1. Then each degree 2 vertex in the $K_{1,1,3}$ is either incident to a double edge joining it to a third vertex, or to

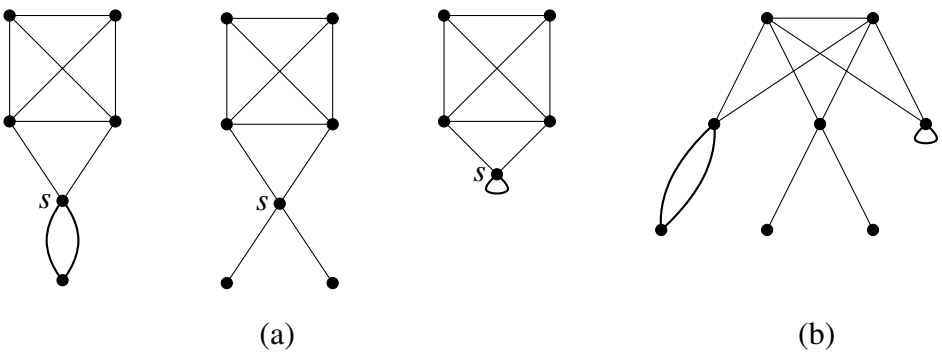

Figure 9: (a) The possibilities in Case 3.a when G_0 contains an induced balanced copy of K_4 and (b) the possibilities when G_0 contains an induced balanced copy of $K_{1,1,3}$.

two single edges joining it to a third and fourth vertex, or to a loop. (These three possibilities are illustrated in Figure 9(b).) If there is a vertex of the first or the second type, then we may apply Lemma 3.13 or Lemma 3.15, respectively, to show that there exists an admissible reverse move. So suppose each of the degree 2 vertices of $K_{1,1,3}$ is incident to a loop. Since $K_{1,1,3}$ with three loops is 4-regular, it must be equal to G_0. We show there is a contraction using Lemma 3.17. Let r and s be the vertices corresponding to the two partite sets of $K_{1,1,3}$ of size 1, and let t be a vertex incident to a loop. We will contract the edge $(rt, 1)$. Lemma 3.3 implies that (i) fails and again (ii) and (iv) clearly fail. To see that (iii) fails note simply that every subgraph H with no loops satisfies $f(H) \geq 3$.

Case 3.b. There exists a vertex not contained in a balanced subgraph isomorphic to K_4 or to $K_{1,1,3}$.

We consider each of the possibilities for the neighbourhood of a given vertex v in G_0 satisfying the hypotheses of this case.

Suppose first that v has exactly one neighbour. Since $G_0 \neq 2K_2^1$ and G_0 is connected, Lemma 3.11 implies that v is admissible. Suppose next that v has either exactly three or exactly four neighbours. Then Lemmas 3.13 and 3.15 imply that v is admissible.

So we suppose that v has exactly two neighbours. Suppose first that there is no loop on v. Let $N(v) = \{x, y\}$. Lemma 3.12 implies that if v is not admissible then there is a loop at x or at y, say at x. Since G_0 is 4-regular we see that x satisfies the condition of Lemma 3.11 and hence there is an admissible reverse move.

So suppose that v has exactly two neighbours x, y and a loop on v. By switching, we may assume that $(vx, 1)$ and $(vy, 1)$ are the non-loop edges incident to v. We consider a possible H2d-reduction at v. If every vertex of G_0 is incident with a loop and has two neighbours then it is easy to see that G_0 is precisely a cycle with one loop on each vertex. If the cycle has length 3 then G_0 is a base graph. If the cycle has length at least 4 then it is easy to see that there is no balanced subgraph H_0 of $G_0 - v$ containing x, y with $f(H_0) = 2$ such that all walks from x to y have gain 1. Hence v is admissible.

So we may suppose there is some vertex in G_0 which is not incident to a loop. By 4-regularity of G_0 and by relabelling if necessary, we may suppose that x is such a vertex. Since x is adjacent to v with exactly one edge and x does not have a loop, x may have either 3 or 4 neighbours.

If x has 3 neighbours then 4-regularity implies that x is not contained in a K_4. Hence Lemma 3.13 implies that x, or one of its neighbours, is admissible. So we suppose that x has 4 distinct neighbours. If the edge $(xy, 1)$ exists, then we may contract the edge vx using Lemma 3.17. (The structure of G_0 makes all four conditions in that lemma easy to rule out.) So suppose $(xy, 1)$ does not exist. If $(xy, -1)$ exists then we claim that x is admissible. If not, then it follows from Lemma 3.15 that x must be in a balanced K_4. This K_4 must consist of the vertices x, y, a, b. By switching we may assume that the three edges joining x with a, b and y all have gain -1 and that the remaining edges of the K_4 have gain 1. Note that every path from x to y within the K_4 has gain -1. This implies that there is no balanced subgraph H_{xy} of $G_0 - v$ containing x and y with $f(H_{xy}) = 2$ such that each path from x to y in H_{xy} has gain 1. Since there is clearly also no subgraph H_0 of $G_0 - v$ containing x and y with $f(H_0) = 0$, this implies that there is a H2d-reduction at v, contradicting our assumption that v is non-admissible.

So we may suppose that x is not adjacent to y. By Lemma 3.15 x is either in a balanced $K_{1,1,3}$, which gives an immediate contradiction, or x is in a balanced K_4. In the latter case, since we are not in Case 3.a, either the K_4 induces exactly one additional edge and we can apply a reverse vertex-to-K_4 move by Lemma 3.16 (since the additional edge cannot be a loop, by 4-regularity of G_0), or the K_4 is contained in a copy of $K_5 - e$.

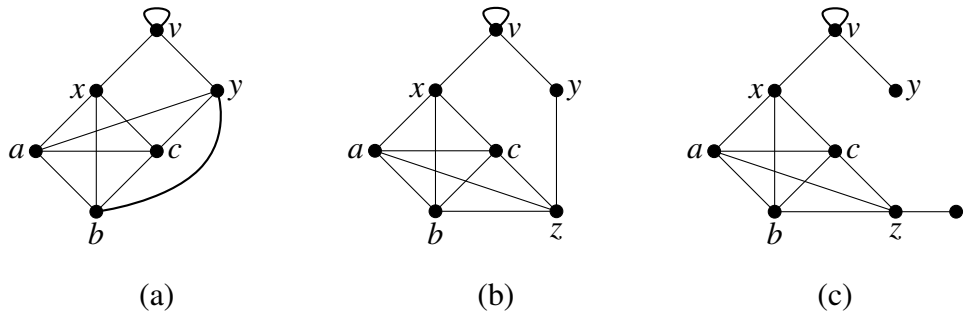

Figure 10: Possible structures of G when v has two neighbours and a loop and one of its neighbours has 4 distinct neighbours.

Let the vertices of the K_4 be a, b, c, x and suppose all edges of the K_4 have gain 1. If the final vertex of the $K_5 - e$ is y then $V(G_0) = \{v, x, y, a, b, c\}$ (see Figure 10(a)) and we claim that y is admissible. To see this, note that exactly one or exactly two of the edges joining y with a, b, c have gain -1, for otherwise K_4 together with the three edges incident with y would violate $(2, 2)$-sparsity. Hence in both cases, y is admissible by Lemma 3.15.

So we may assume that the final vertex of the $K_5 - e$ is $z \neq y$. Suppose first that z is adjacent to y (see Figure 10(b)). Then we claim that y is still admissible. Clearly y is adjacent to v and z. If y is also incident to a loop, then $V(G_0) = \{v, x, y, a, b, c, z\}$ and it is easy to see that a H2d-reduction is possible at y. If y is incident to two parallel edges joining it with a vertex $u \neq v, z$ then y, or one of its neighbours, is admissible by Lemma 3.13 and the 4-regularity of G_0. If y is incident to two

single edges joining it with vertices u and w that are distinct from v and z, then y is admissible by Lemma 3.15 and the 4-regularity of G_0. So we may suppose that z is not adjacent to y (see Figure 10(c)). Then z is admissible by Lemma 3.15 and the 4-regularity of G_0.

Case 3.c. Every vertex is contained in a balanced subgraph isomorphic to K_4 plus either one or two additional edges or $K_5 - e$.

If G_0 contains a K_4^{++} as a subgraph then $G_0 = K_4^{++}$, by 4-regularity of G_0. In this case G_0 is the base graph depicted in Figure 1(h). So we may assume that every vertex of G_0 is contained in a subgraph isomorphic to K_4^+ (which cannot have a loop) or $K_5 - e$. Since G_0 is 4-regular any pair of copies of K_4^+ or $K_5 - e$ are vertex disjoint, and each copy has exactly two edges incident to it. If G_0 contains a copy of K_4^+ then we can apply a reverse vertex-to-K_4 move by Lemma 3.16. (Note that since every vertex is in a K_4^+ or a $K_5 - e$, and G_0 is 4-regular, there cannot exist a vertex outside of the K_4^+ that is adjacent to two of the vertices of the K_4^+.) Hence we may suppose that G_0 contains no copies of K_4^+. Thus we may assume that we have a copy of $K_5 - e$. If $G_0 = K_5$ then it is elementary to apply Lemma 3.15 to find an admissible reduction. Hence G_0 does not contain a copy of K_5. Suppose $e = xy$ and note that if x (resp. y) is not admissible, then by Lemma 3.15 x (resp. y) must be contained in a balanced K_4. However, if x and y are both contained in balanced K_4's, then the $K_5 - e$ is the union of two balanced subgraphs and is hence balanced by Lemma 3.2, contradicting $(2, 2)$-sparsity. □

4. Application to C_2-symmetric frameworks in the ℓ^1 and ℓ^∞-plane

Let $\|\cdot\|_{\mathcal{P}}$ be a norm on $\mathbb{R}^2$ with the property that the closed unit ball $\mathcal{P} = \{x \in \mathbb{R}^2 : \|x\|_{\mathcal{P}} \leq 1\}$ is a quadrilateral (eg. the ℓ^1 or ℓ^∞ norms). We refer to bar-joint frameworks in this context as *grid-like*. In this section, the results of the previous sections are combined to obtain geometric and combinatorial characterisations of χ-symmetric isostaticity and infinitesimal rigidity for C_2-symmetric grid-like frameworks.

4.1. Framework colourings

Let (G, p) be a well-positioned grid-like bar-joint framework and let $F \in \{\pm F_1, \pm F_2\}$ be one of the four facets of the quadrilateral $\mathcal{P}$. An edge $vw \in E$ is said to have *framework colour* F (equivalently, $-F$) if either $p_v - p_w$ or $p_w - p_v$ lies in the cone $\{x \in \mathbb{R}^2 : \frac{x}{\|x\|_{\mathcal{P}}} \in F\}$. Recall that, since (G, p) is well-positioned, each edge of G has exactly one framework colour (see [5]). Denote by G_F the *monochrome subgraph* of G spanned by edges with framework colour F.

For each facet F there exists a unique extreme point $\hat{F}$ of the polar set $\mathcal{P}^\triangle = \{y \in \mathbb{R}^2 : x \cdot y \leq 1, \ \forall\, x \in \mathcal{P}\}$ such that $F = \{x \in \mathcal{P} : x \cdot \hat{F} = 1\}$. Define a linear functional $\varphi_F : \mathbb{R}^2 \to \mathbb{R}$ by setting $\varphi(x) = x \cdot \hat{F}$, for all $x \in \mathbb{R}^2$. If (G, p) is well-positioned and $vw \in G_F$ then it can be shown (see [5]) that the linear functional $\varphi_{v,w}$ described in Lemma 2.1 satisfies $\varphi_{v,w} = \varphi_F$.

If $\mathcal{G} = (G, p, \theta, \tau)$ is a C_2-symmetric grid-like bar-joint framework, then each edge $e \in E$ shares the same framework colour as its image $-e$. By assigning this common framework colour to the edge orbit $[e] = \{e, -e\}$ we induce a framework colouring on the edges of the quotient graph G_0. Denote by $G_{F,0}$ the *monochrome subgraph* of G_0 spanned by edges $[e]$ with framework colour F.

Example 4.1. Consider the ℓ^∞ plane. The unit ball $\mathcal{P} = \{x \in \mathbb{R}^2 : \|x\|_\infty \leq 1\}$ has four facets: $F_1 = \{(x_1, x_2) \in \mathcal{P} : x_1 = 1\}$, $F_2 = \{(x_1, x_2) \in \mathcal{P} : x_2 = 1\}$ and their negatives. The polar set of $\mathcal{P}$ is the ℓ^1 unit ball $\mathcal{P}^\triangle = \{x \in \mathbb{R}^2 : \|x\|_1 \leq 1\}$, and the extreme points of the polar set are $\hat{F}_1 = (1, 0)$, $\hat{F}_2 = (0, 1)$ and their negatives. Figure 1 illustrates several examples of framework colourings for C_2-symmetric bar-joint frameworks in the ℓ^∞-plane together with the induced framework colourings on their $\mathbb{Z}_2$-gain graphs.

A *map graph* is a graph in which every connected component contains exactly one cycle. An *unbalanced map graph* is a $\mathbb{Z}_2$-gain graph (H, ψ) such that H is a map graph and every cycle is unbalanced.

Theorem 4.2. *Let $\mathcal{G} = (G, p, \theta, \tau)$ be a well-positioned and C_2-symmetric bar-joint framework in $(\mathbb{R}^2, \|\cdot\|_\mathcal{P})$.*

(A) The following statements are equivalent.

(i) $\mathcal{G}$ is χ_0-symmetrically isostatic.

(ii) $G_{F_1,0}$ and $G_{F_2,0}$ are edge-disjoint spanning unbalanced map graphs in G_0.

(B) The following statements are equivalent.

(i) $\mathcal{G}$ is χ_1-symmetrically isostatic.

(ii) $G_{F_1,0}$ and $G_{F_2,0}$ are edge-disjoint spanning trees in G_0.

(C) The following statements are equivalent.

(i) $\mathcal{G}$ is infinitesimally rigid.

(ii) $G_{F_1,0}$ and $G_{F_2,0}$ both contain connected spanning unbalanced map graphs.

Proof. (A) $(i) \Rightarrow (ii)$ Suppose there exists a vertex $[v_0] \in V_0 \setminus V(G_{F_1,0})$. Let $\tilde{v}_0$ be the representative vertex for $[v_0]$ in G. Choose a non-zero vector $x \in \ker \varphi_{F_2}$ and for all $v \in V(G)$ define,

$$u_v = \begin{cases} x & \text{if } v = \tilde{v}_0, \\ -x & \text{if } v = -\tilde{v}_0, \\ 0 & \text{otherwise.} \end{cases}$$

Then u is a non-trivial χ_0-symmetric infinitesimal flex for (G, p). This is a contradiction since (G, p) is χ_0-symmetrically isostatic. Thus every vertex of G_0 must be incident to an edge of $G_{F_1,0}$. By a similar argument every vertex of G_0 must be incident to an edge of $G_{F_2,0}$.

Suppose $G_{F_1,0}$ has a connected component H_0 which is a balanced subgraph of G_0. Then, by Lemma 2.17, we may assume that each edge of H_0 has gain 1. Thus if H is the covering graph for H_0, then there is no edge $vw \in E(H)$ with $v \in \tilde{V}_0$ and $w \notin \tilde{V}_0$. (Recall Section 2.8 for the definition of $\tilde{V}_0$). Choose a non-zero vector $x \in \ker \varphi_{F_2}$ and for all $v \in V(G)$ define,

$$u_v = \begin{cases} x & \text{if } [v] \in V(H_0) \text{ and } v \in \tilde{V}_0, \\ -x & \text{if } [v] \in V(H_0) \text{ and } v \notin \tilde{V}_0, \\ 0 & \text{otherwise.} \end{cases}$$

Then u is a non-trivial χ_0-symmetric infinitesimal flex for (G, p). This is a contradiction and so every connected component of $G_{F_1,0}$ must be an unbalanced subgraph of G_0. Similarly, each connected component of $G_{F_2,0}$ is an unbalanced subgraph of G_0.

By Corollary 2.20, we have $|E_0| = 2|V_0|$. Note that each connected component of $G_{F_1,0}$ must contain a cycle (since it is unbalanced) and so if $G_{F_1,0}$ has n connected components, $H_1, H_2, \ldots, H_n$ say, then $|E(H_j)| \geq |V(H_j)|$ for each j and,

$$|E(G_{F_1,0})| = \sum_{j=1}^{n} |E(H_j)| \geq \sum_{j=1}^{n} |V(H_j)| = |V_0|.$$

Similarly, $|E(G_{F_2,0})| \geq |V_0|$. Now $|E(G_{F_1,0})| + |E(G_{F_2,0})| = |E_0| = 2|V_0|$ and so $|E(G_{F_1,0})| = |V_0| = |E(G_{F_2,0})|$. It follows that $|E(H_j)| = |V(H_j)|$ for each j and so the connected components of $G_{F_1,0}$ each contain exactly one cycle. By the same argument, the connected components of $G_{F_2,0}$ each contain exactly one cycle. Thus $G_{F_1,0}$ and $G_{F_2,0}$ are both unbalanced spanning mapping graphs in G_0.

$(ii) \Rightarrow (i)$ Suppose (ii) holds and let u be a χ_0-symmetric infinitesimal flex of (G, p). Then $u_{-v} = -u_v$ for all $v \in V$. Let $v_0 \in V$ and let H_0^1 and H_0^2 be the connected components of $G_{F_1,0}$ and $G_{F_2,0}$ respectively which contain $[v_0] \in V_0$. Since H_0^i contains a unique unbalanced cycle, there exists a path in G_{F_i} from v_0 to $-v_0$. It follows that $u_{v_0} - u_{-v_0} \in \cap_{i=1,2} \ker \varphi_{F_i} = \{0\}$ and so $u_{v_0} = u_{-v_0} = -u_{v_0}$. Thus $u_{v_0} = 0$. Applying this argument to all $v \in V$, we have $u = 0$ and so (G, p, θ, τ) is χ_0-symmetrically infinitesimally rigid. Note that $|E_0| = 2|V_0|$ and so $\mathcal{G}$ is also χ_0-symmetrically isostatic.

(B) $(i) \Rightarrow (ii)$ Suppose there exists a vertex $[v_0] \in V_0 \setminus V(G_{F_1,0})$. Choose a non-zero vector $x \in \ker \varphi_{F_2}$. For all $v \in V$ define,

$$u_v = \begin{cases} x & \text{if } [v] = [v_0], \\ 0 & \text{otherwise.} \end{cases}$$

Then u is a non-trivial χ_1-symmetric infinitesimal flex for (G, p). This is a contradiction and so $G_{F_1,0}$ is a spanning subgraph of G_0. Similarly, $G_{F_2,0}$ is a spanning subgraph of G_0.

Suppose $G_{F_1,0}$ is not connected, and let H_0 be a connected component of $G_{F_1,0}$. Choose a non-zero vector $x \in \ker \varphi_{F_2}$ and for all $v \in V$ define,

$$u_v = \begin{cases} x & \text{if } [v] \in V(H_0), \\ 0 & \text{otherwise.} \end{cases}$$

Then u is a non-trivial χ_1-symmetric infinitesimal flex for (G, p), which is a contradiction. Thus $G_{F_1,0}$ is a connected spanning subgraph of G_0. Similarly, $G_{F_2,0}$ is a connected spanning subgraph of G_0. By Corollary 2.20, we have $|E_0| = 2|V_0| - 2$. Note that $|E(G_{F_1,0})| \geq |V_0| - 1$ and $|E(G_{F_2,0})| \geq |V_0| - 1$ and so $G_{F_1,0}$ and $G_{F_2,0}$ are both spanning trees in G_0.

$(ii) \Rightarrow (i)$ Suppose (ii) holds and let u be a χ_1-symmetric infinitesimal flex for $\mathcal{G}$. Then $u_{-v} = u_v$ for all $v \in V$. Fix $v, w \in V$. Since $G_{F_1,0}$ is a spanning tree in G_0, there exists a path in $G_{F_1,0}$ from $[v]$ to $[w]$. Thus there either exists a path P in G_{F_1} from v to w or there exists a path P in G_{F_1} from v to $-w$. In the former case it follows directly that $u_v - u_w \in \ker \varphi_{F_1}$ while in the latter case

it follows that $u_v - u_w = u_v - u_{-w} \in \ker \varphi_{F_1}$. Similarly, $u_v - u_w \in \ker \varphi_{F_2}$ and so $u_v = u_w$ for all $v, w \in V$. Thus u is a trivial infinitesimal flex and so $\mathcal{G}$ is χ_1-symmetrically infinitesimally rigid. Since $|E_0| = 2|V_0| - 2$, $\mathcal{G}$ is also χ_1-symmetrically isostatic.

(C) $(i) \Rightarrow (ii)$ If (G, p) is infinitesimally rigid then it is both χ_0 and χ_1-symmetrically infinitesimally rigid. By removing edge orbits from G we arrive at a spanning subgraph A such that (A, p) is χ_0-symmetrically isostatic. By (A), $A_{F_1,0}$ and $A_{F_2,0}$ are unbalanced spanning map graphs in G_0 and so each contains an unbalanced cycle. Similarly, (G, p) contains a spanning subgraph B such that (B, p) is χ_1-symmetrically isostatic. By (B), $B_{F_1,0}$ and $B_{F_2,0}$ are spanning trees in G_0. Since $B_{F_i,0}$ is a spanning tree for $i = 1, 2$, there exists a set of edges in $B_{F_i,0}$ which, when added to $A_{F_i,0}$, form a connected unbalanced spanning map graph $H_{F_i,0}$ after removing edges to reduce the number of cycles to one. This gives the result.

$(ii) \Rightarrow (i)$ Suppose (ii) holds. Then (G_0, ψ) contains a spanning subgraph H_0 such that the induced monochrome subgraphs $H_{F_1,0}$ and $H_{F_2,0}$ are edge-disjoint connected unbalanced spanning map graphs. Let H be the covering graph for H_0. By (A), the C_2-symmetric subframework (H, p) is χ_0-symmetrically infinitesimally rigid. Similarly, note that $H_{F_1,0}$ and $H_{F_2,0}$ both contain spanning trees in H_0 and so by (B), (H, p) is χ_1-symmetrically infinitesimally rigid. It follows that (H, p), and hence also (G, p), is infinitesimally rigid. □

4.2. Existence of rigid grid-like placements with half-turn symmetry

Recall from Corollary 2.20 that if $\mathcal{G} = (G, p, \theta, \tau)$ is a well-positioned, C_2-symmetric and χ_0-symmetrically isostatic bar-joint framework in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$, where $\mathcal{P}$ is a quadrilateral, then the gain graph (G_0, ψ) for (G, θ) is $(2, 2, 0)$-gain-tight. By Theorem 3.20, (G_0, ψ) is $(2, 2, 0)$-gain-tight if it can be generated from vertex-disjoint copies of graphs in $\mathcal{B}$ by applying H1, H2, H3, vertex-to-K_4 and vertex splitting moves. We now show that if there exists such a recursive construction sequence, then there exists a half-turn symmetric realisation of G that is well-positioned and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$. Overall, this yields the following main combinatorial result for χ_0-symmetrically isostatic frameworks with half-turn symmetry in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$.

Theorem 4.3. *Let $\|\cdot\|_{\mathcal{P}}$ be a norm on $\mathbb{R}^2$ for which $\mathcal{P}$ is a quadrilateral, and let (G, θ) be a $\mathbb{Z}_2$-symmetric graph. Further, let (G_0, ψ) be the gain graph for (G, θ). The following are equivalent.*

(i) There exists a C_2-symmetric realisation $\mathcal{G} = (G, p, \theta, \tau)$ which is well-positioned and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$;

(ii) (G_0, ψ) is $(2, 2, 0)$-gain tight;

(iii) (G_0, ψ) can be constructed from disjoint copies of base graphs in Figure 1 by a sequence of H1a,b,c moves, H2a,b,c,d,e moves, H3a,b,c,d moves, vertex-to-K_4 moves, and vertex splitting moves.

To show that (iii) implies (i), we rely on Theorem 4.2(A). We split the proof into a number of geometric lemmas. In these lemmas, we will use the notation of Section 2.8 and write $[v]$ and $[e]$ for a vertex and an edge of the gain graph (G_0, ψ) for a $\mathbb{Z}_2$-symmetric graph (G, θ), respectively. Moreover, we let $\tilde{V}_0 = \{\tilde{v}_1, \ldots, \tilde{v}_n\}$ be a choice of representatives for the vertex orbits of (G, θ).

Lemma 4.4. *Let (G_0, ψ) and (G'_0, ψ') be the gain graphs of the $\mathbb{Z}_2$-symmetric graphs (G, θ) and (G', θ'), respectively and suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a H1a, H1b or H1c move. If for (G', θ') there exists a realisation $\mathcal{G} = (G', p', \theta', \tau)$ which is well-positioned, C_2-symmetric and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$, then the same is true for (G, θ).*

Proof. By Theorem 4.2(A), there exists a well-positioned C_2-symmetric realisation p' of (G', θ') in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$ so that the induced monochrome subgraphs $G'_{F_1,0}$ and $G'_{F_2,0}$ of (G'_0, ψ') are both spanning unbalanced map graphs. By Theorem 4.2(A), it now suffices to show that the vertex of $G_0 \setminus G'_0$ can be placed in such a way that the corresponding framework (G, p, θ, τ) is C_2-symmetric and well-positioned, and the induced monochrome subgraphs $G_{F_1,0}$ and $G_{F_2,0}$ are both spanning unbalanced map graphs in (G_0, ψ).

We fix two points x_1 and x_2 in the relative interiors of F_1 and F_2 respectively.

Suppose first that (G_0, ψ) is obtained from (G'_0, ψ') by a H1a move, where $[v] \in G_0 \setminus G'_0$ is adjacent to the vertices $[v_1]$ and $[v_2]$ of G'_0 with respective gains γ_1 and γ_2. Set $p_w = p'_w$ for all vertices w of G with $[w] \neq [v]$. Let $a \in \mathbb{R}^2$ be the point of intersection of the lines $L_1 = \{\tau(\gamma_1)p_{\tilde{v}_1} + tx_1 : t \in \mathbb{R}\}$ and $L_2 = \{\tau(\gamma_2)p_{\tilde{v}_2} + tx_2 : t \in \mathbb{R}\}$ and let $B(a, r)$ be an open ball with centre a and radius $r > 0$. Choose $p_{\tilde{v}}$ to be any point in $B(a, r)$ which is distinct from $\{p_w : w \in V(G')\}$ and which is not fixed by $\tau(-1)$. Set $p_{-\tilde{v}} = \tau(-1)p_{\tilde{v}}$. Then (G, p, θ, τ) is a C_2-symmetric bar-joint framework and, by applying a small perturbation to $p_{\tilde{v}}$ if necessary, we may assume that (G, p, θ, τ) is well-positioned. If r is sufficiently small then the induced framework colours for $[v][v_1]$ and $[v][v_2]$ are $[F_1]$ and $[F_2]$ respectively. Thus, the induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = G'_{F_1,0} \cup \{[v][v_1]\}$ and $G_{F_2,0} = G'_{F_2,0} \cup \{[v][v_2]\}$. Clearly, $G_{F_1,0}$ and $G_{F_2,0}$ are spanning unbalanced map graphs of (G_0, ψ).

If (G_0, ψ) is obtained from (G'_0, ψ) by a H1b move, then the proof is completely analogous to the proof above.

Suppose (G_0, ψ) is obtained from (G'_0, ψ') by a H1c move, where $[v] \in G_0 \setminus G'_0$ is incident to the unbalanced loop $[e]$ and adjacent to the vertex $[z]$ of (G'_0, ψ') with gain γ. Let $a \in \mathbb{R}^2$ be the point of intersection of the lines $L_1 = \{\tau(\gamma)p_{\tilde{z}} + tx_2 : t \in \mathbb{R}\}$ and $L_2 = \{tx_1 : t \in \mathbb{R}\}$ and let $B(a, r)$ be an open ball with centre a and radius $r > 0$. Choose $p_{\tilde{v}}$ to be any point in $B(a, r)$ which is distinct from $\{p_w : w \in V(G')\}$ and which is not fixed by $\tau(-1)$. Set $p_{-\tilde{v}} = \tau(-1)p_{\tilde{v}}$. Then, by applying a small perturbation to $p_{\tilde{v}}$ if necessary, (G, p, θ, τ) is well-positioned and C_2-symmetric. Moreover, if r is sufficiently small, then the induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = G'_{F_1,0} \cup \{[e]\}$ and $G_{F_2,0} = G'_{F_2,0} \cup \{[v][z]\}$. Clearly, $G_{F_1,0}$ and $G_{F_2,0}$ are unbalanced spanning map graphs of (G_0, ψ). □

Lemma 4.5. *Let (G_0, ψ) and (G'_0, ψ') be the gain graphs of the $\mathbb{Z}_2$-symmetric graphs (G, θ) and (G', θ'), respectively and suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a H2a, H2b, H2c, H2d, or H2e move. If for (G', θ') there exists a realisation $\mathcal{G} = (G', p', \theta', \tau)$ which is well-positioned, C_2-symmetric and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$, then the same is true for (G, θ).*

Proof. By Theorem 4.2(A), there exists a well-positioned C_2-symmetric realisation p' of (G', θ') in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$ so that the induced monochrome subgraphs $G'_{F_1,0}$ and $G'_{F_2,0}$ of (G'_0, ψ') are both spanning unbalanced map graphs. By Theorem 4.2(A), it now suffices to show that the vertex of $G_0 \setminus G'_0$ can

be placed in such a way that the corresponding framework (G, p, θ, τ) is C_2-symmetric and well-positioned, and the induced monochrome subgraphs $G_{F_1,0}$ and $G_{F_2,0}$ are both spanning unbalanced map graphs in (G_0, ψ).

We fix two points x_1 and x_2 in the relative interiors of F_1 and F_2 respectively.

Suppose first that (G_0, ψ) is obtained from (G'_0, ψ') by a H2a move where $[v] \in G_0 \setminus G'_0$ subdivides the edge $[e] = [v_1][v_2]$ into the edges $[e_1] = [v][v_1]$ and $[e_2] = [v][v_2]$ with respective gains γ_1 and γ_2, and $[v]$ is also incident to the edge $[e_3]$ with end-vertex $[z]$ and gain γ_3. Without loss of generality we may assume that $[e] \in G'_{F_1,0}$. Let $a \in \mathbb{R}^2$ be the point of intersection of the line L_1 which passes through the points $\tau(\gamma_1)p_{\tilde{v}_1}$ and $\tau(\gamma_2)p_{\tilde{v}_2}$, and the line $L_2 = \{\tau(\gamma_3)p_{\tilde{z}} + tx_2 : t \in \mathbb{R}\}$. Let $B(a, r)$ be the open ball with centre a and radius $r > 0$ and choose $p_{\tilde{v}}$ to be a point in $B(a, r)$ which is distinct from $\{p_w : w \in G'\}$ and which is not fixed by $\tau(-1)$. Set $p_{-\tilde{v}} = \tau(-1)p_{\tilde{v}}$. Then (G, p, θ, τ) is C_2-symmetric and, by applying a small perturbation to $p_{\tilde{v}}$ if necessary, we may assume it is well-positioned. If r is sufficiently small then $[e_1]$ and $[e_2]$ have induced framework colour $[F_1]$ and $[e_3]$ has framework colour $[F_2]$. The induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = (G'_{F_1,0} \setminus \{[e]\}) \cup \{[e_1], [e_2]\}$ and $G_{F_2,0} = G'_{F_2,0} \cup \{[e_3]\}$. Clearly, $G_{F_1,0}$ and $G_{F_2,0}$ are spanning unbalanced map graphs of (G_0, ψ).

The cases where (G_0, ψ) is obtained from (G'_0, ψ') by a H2b or a H2c move may be proved completely analogously to the case above for the H2a move.

Next, we suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a H2d move where $[v] \in G_0 \setminus G'_0$ subdivides the edge $[e] = [v_1][v_2]$ into the edges $[e_1] = [v][v_1]$ and $[e_2] = [v][v_2]$ with respective gains γ_1 and γ_2, and $[v]$ is also incident to the unbalanced loop $[e_3]$. Without loss of generality we may assume that $[e] \in G'_{F_1,0}$. Let $a \in \mathbb{R}^2$ be the point of intersection of the line L_1 which passes through the points $\tau(\gamma_1)p_{\tilde{v}_1}$ and $\tau(\gamma_2)p_{\tilde{v}_2}$ and the line $L_2 = \{tx_2 : t \in \mathbb{R}\}$, and let $B(a, r)$ be an open ball with centre a and radius $r > 0$. (Note that a could possibly be the centre of the rotation $\tau(-1)$, i.e. the origin.) Choose $p_{\tilde{v}}$ to be a point in $B(a, r)$ which is distinct from $\{p_w : w \in G'\}$ and which is not fixed by $\tau(-1)$. Set $p_{-\tilde{v}} = \tau(-1)p_{\tilde{v}}$. Then (G, p, θ, τ) is C_2-symmetric and if r is sufficiently small then $[e_1]$ and $[e_2]$ have induced framework colour $[F_1]$. Moreover, by applying a perturbation to $p_{\tilde{v}}$ within $B(a, r)$ if necessary, we may assume that $[e_3]$ has framework colour $[F_2]$ and that (G, p, θ, τ) is well-positioned. The induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = (G'_{F_1,0} \setminus \{[e]\}) \cup \{[e_1], [e_2]\}$ and $G_{F_2,0} = G'_{F_2,0} \cup \{[e_3]\}$. Clearly, $G_{F_1,0}$ and $G_{F_2,0}$ are spanning unbalanced map graphs of (G_0, ψ).

The case where (G_0, ψ) is obtained from (G'_0, ψ') by a H2e move may be proved completely analogously to the case above for the H2d move. Here $B(a, r)$ is in fact centred at the origin. □

Lemma 4.6. *Let (G_0, ψ) and (G'_0, ψ') be the gain graphs of the $\mathbb{Z}_2$-symmetric graphs (G, θ) and (G', θ'), respectively and suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a H3a, H3b, H3c, or H3d move. If for (G', θ') there exists a realisation $\mathcal{G} = (G', p', \theta', \tau)$ which is well-positioned, C_2-symmetric and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$, then the same is true for (G, θ).*

Proof. By Theorem 4.2(A), there exists a well-positioned C_2-symmetric realisation p' of (G', θ') in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$ so that the induced monochrome subgraphs $G'_{F_1,0}$ and $G'_{F_2,0}$ of (G'_0, ψ') are both spanning unbalanced map graphs. By Theorem 4.2(A), it now suffices to show that the vertex of $G_0 \setminus G'_0$ can be placed in such a way that the corresponding framework (G, p, θ, τ) is C_2-symmetric and well-

positioned, and the induced monochrome subgraphs $G_{F_1,0}$ and $G_{F_2,0}$ are both spanning unbalanced map graphs in (G_0, ψ).

First we suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a H3a move where $[v] \in G_0 \setminus G'_0$ subdivides the edge $[e] = [v_1][v_2]$ into the edges $[e_1] = [v][v_1]$ and $[e_2] = [v][v_2]$, and the edge $[f] = [v_3][v_4]$ into the edges $[f_1] = [v][v_3]$ and $[f_2] = [v][v_4]$. By switching $[v_1]$ and $[v_3]$ if necessary, we may assume without loss of generality that $[e]$ and $[f]$ have both gain 1. The edges $[e_1], [e_2], [f_1], [f_2]$ will then also be assigned gain 1. (The proof for the case where they are all assigned gain -1 is analogous.) We distinguish two cases.

Case A: $[e]$ and $[f]$ belong to different induced monochrome subgraphs of G'_0, say $[e] \in G'_{F_1,0}$ and $[f] \in G'_{F_2,0}$. Let $a \in \mathbb{R}^2$ be the point of intersection of the line L_1 which passes through the points $p_{\tilde{v}_1}$ and $p_{\tilde{v}_2}$, and the line L_2 which passes through the points $p_{\tilde{v}_3}$ and $p_{\tilde{v}_4}$. Let $B(a, r)$ be an open ball with centre a and radius $r > 0$, and choose $p_{\tilde{v}}$ to be a point in $B(a, r)$ which is distinct from $\{p_w : w \in G'\}$ and which is not fixed by $\tau(-1)$. Set $p_{-\tilde{v}} = \tau(-1)p_{\tilde{v}}$. Then (G, p, θ, τ) is C_2-symmetric and, by applying a small perturbation to $p_{\tilde{v}}$ if necessary, we may assume it is well-positioned. If r is sufficiently small then $[e_1]$ and $[e_2]$ have induced framework colour $[F_1]$, and $[f_1]$ and $[f_2]$ have framework colour $[F_2]$. The induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = (G'_{F_1,0} \setminus \{[e]\}) \cup \{[e_1], [e_2]\}$ and $G_{F_2,0} = (G'_{F_2,0} \setminus \{[f]\}) \cup \{[f_1], [f_2]\}$. Clearly, $G_{F_1,0}$ and $G_{F_2,0}$ are spanning unbalanced map graphs of (G_0, ψ).

Case B: $[e]$ and $[f]$ belong to the same induced monochrome subgraph of G'_0, say $[e], [f] \in G'_{F_1,0}$. We need the following claim.

Claim 4.7. Let p_1, p_2, p_3, p_4 be four distinct points in $\mathbb{R}^2$ such that the line segments p_1p_2 and p_3p_4 both have framework colour $[F_1]$. Let $i \in \{1, 2, 3, 4\}$. Then there exists an open set N in $\mathbb{R}^2$ such that for every point $p_v \in N$, the line segment p_vp_i has framework colour $[F_2]$ and the three line segments p_vp_j with $j \in \{1, 2, 3, 4\}$, $j \neq i$, have framework colour $[F_1]$.

Proof. Without loss of generality we may assume that p_1 lies to the left of p_2, and p_3 lies to the left of p_4 (see Figure 11 for an illustration). Moreover we may assume that $i = 4$.

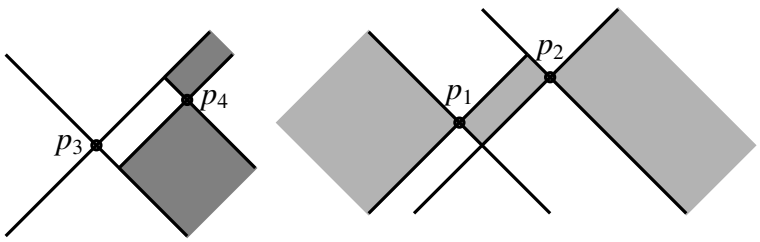

Figure 11: Illustration of the proof of Claim 4.7.

We need to find an open set N which lies within the two shaded areas in Figure 11. Note that the shaded area on the left hand side of Figure 11 is connected, and unbounded from below and above. The shaded area on the right hand side of Figure 11 is also connected, and unbounded from the left and right. Since p_1, p_2, p_3, p_4 are distinct points, the shaded areas will always have a nontrivial intersection, and we may choose N within that intersection. □

Suppose first that $[e]$ and $[f]$ lie on a common (unbalanced) cycle in $G'_{F_1,0}$. Without loss of generality, we may assume that this cycle consists of a path P_1 from $[v_2]$ to $[v_3]$ with an odd number of edges with gain -1, and a path P_2 from $[v_1]$ to $[v_4]$ with an even number of edges with gain -1. Then we choose $p_{\tilde{v}}$ to be a point which is distinct from $\{p_w : w \in G'\}$, not fixed by $\tau(-1)$, and such that (G, p, θ, τ) is well-positioned and in (G_0, ψ) the edges $[e_1], [e_2], [f_1]$ have framework colour $[F_1]$, and $[f_2]$ has framework colour $[F_2]$. Such a position for $p_{\tilde{v}}$ exists by Claim 4.7. The induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = (G'_{F_1,0} \backslash \{[e], [f]\}) \cup \{[e_1], [e_2], [f_1]\}$ and $G_{F_2,0} = (G'_{F_2,0}) \cup \{[f_2]\}$. Clearly, $G_{F_2,0}$ is a spanning unbalanced map graph of (G_0, ψ). As for $G_{F_1,0}$, note that the removal of $[e]$ and $[f]$ from $G'_{F_1,0}$ breaks the connected component of $G'_{F_1,0}$ containing $[e]$ and $[f]$ into the two disjoint trees. By adding the vertex $[v]$ and the edges $[e_1], [e_2], [f_1]$, these two trees are reconnected and a single unbalanced cycle (consisting of P_1, $[e_2]$ and $[f_1]$) is created in this connected component of $G_{F_1,0}$.

If $[e]$ and $[f]$ do not lie on a common cycle in $G'_{F_1,0}$, but they are still in the same connected component K' of $G'_{F_1,0}$, then we may proceed as above. However, if either $[e]$ or $[f]$, say $[e]$, lies on the unique cycle C' in K', and without loss of generality there exists a path in K' from a vertex in C' to $[v_4]$ that does not include $[v_3]$, then we need to choose $p_{\tilde{v}}$ so that $[e_1]$ and $[e_2]$ both have framework colour $[F_1]$, and $[f_1]$ and $[f_2]$ have respective framework colours $[F_1]$ and $[F_2]$. This guarantees that the unbalanced cycle $C = C' \backslash \{[e]\} \cup \{[e_1], [e_2]\}$ in the corresponding component of $G_{F_1,0}$ is unique.

If $[e]$ and $[f]$ lie in different connected components K' and K'' of $G'_{F_1,0}$, then we may again proceed as above. However, care needs to be taken in the case where either $[e]$ or $[f]$, say $[e]$, lies on the unique cycle C' in K', and $[f]$ does not lie on the cycle of K''. In this case we choose $p_{\tilde{v}}$ so that $[e_1]$ and $[e_2]$ have framework colours $[F_1]$ and $[F_2]$, and $[f_1]$ and $[f_2]$ both have framework colour $[F_1]$, so that $G_{F_1,0}$ will not have a connected component with two cycles. Similarly, if neither $[e]$ nor $[f]$ lie on the cycle in their respective connected components, then we need to choose $p_{\tilde{v}}$ so that the three new edges with framework colour $[F_1]$ do not give rise to a connected component of $G_{F_1,0}$ that has two cycles.

Next, we suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a H3b move where $[v] \in G_0 \backslash G'_0$, and the H3b move deletes the edges $[e] = [v_1][v_2]$ and $[f] = [v_1][v_3]$ and adds the edges $[e_1] = [v][v_1]$ and $[e'_1] = [v][v_1]$, and the edges $[e_2] = [v][v_2]$ and $[e_3] = [v][v_3]$. By switching $[v_2]$ and $[v_3]$ if necessary, we may assume that $[e]$ and $[f]$ have gain 1. The edges $[e_1]$ and $[e'_1]$ are assigned the gains 1 and -1, respectively, and the edges $[e_2], [e_3]$ are assigned the gains 1 and -1, respectively. We distinguish two cases.

Case A: $[e]$ and $[f]$ belong to different induced monochrome subgraphs of G'_0, say $[e] \in G'_{F_1,0}$ and $[f] \in G'_{F_2,0}$. Let $a \in \mathbb{R}^2$ be the point of intersection of the line L_1 which passes through the points $p_{\tilde{v}_1}$ and $p_{\tilde{v}_2}$, and the line L_2 which passes through the points $\tau(-1)p_{\tilde{v}_1}$ and $\tau(-1)p_{\tilde{v}_3}$. Let $B(a, r)$ be the open ball with centre a and radius $r > 0$ and choose $p_{\tilde{v}}$ to be a point in $B(a, r)$ which is distinct from $\{p_w : w \in G'\}$ and which is not fixed by $\tau(-1)$. Set $p_{-\tilde{v}} = \tau(-1)p_{\tilde{v}}$. Then (G, p, θ, τ) is C_2-symmetric and, by applying a small perturbation to $p_{\tilde{v}}$ if necessary, we may assume it is well-positioned. If r is sufficiently small then $[e_1]$ and $[e_2]$ have induced framework colour $[F_1]$, and $[e'_1]$ and $[e_3]$ have framework colour $[F_2]$. The induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = (G'_{F_1,0} \backslash \{[e]\}) \cup \{[e_1], [e_2]\}$ and $G_{F_2,0} = (G'_{F_2,0} \backslash \{[f]\}) \cup \{[e'_1], [e_3]\}$. Clearly, $G_{F_1,0}$ and $G_{F_2,0}$

are spanning unbalanced map graphs of (G_0, ψ).

Case B: $[e]$ and $[f]$ belong to the same induced monochrome subgraph of G'_0, say $[e], [f] \in G'_{F_1,0}$.

Suppose first that $[e]$ and $[f]$ lie on a common cycle in $G'_{F_1,0}$. Then we may apply Claim 4.7 to the points $p_{\tilde{v}_1}, \tau(-1)p_{\tilde{v}_1}, p_{\tilde{v}_2}, \tau(-1)p_{\tilde{v}_3}$ to find a position for $p_{\tilde{v}}$ so that it is distinct from $\{p_w : w \in G'\}$, not fixed by $\tau(-1)$, and such that (G, p, θ, τ) is well-positioned and in (G_0, ψ) the edges $[e_1], [e'_1]$ and $[e_2]$ have framework colour $[F_1]$, and $[e_3]$ has framework colour $[F_2]$. The induced monochrome subgraphs of (G_0, ψ) are $G_{F_1,0} = (G'_{F_1,0} \backslash \{[e], [f]\}) \cup \{[e_1], [e'_1], [e_2]\}$ and $G_{F_2,0} = G'_{F_2,0} \cup \{[e_3]\}$. Clearly, $G_{F_1,0}$ and $G_{F_2,0}$ are spanning unbalanced map graphs of (G_0, ψ).

Suppose next that $[e]$ and $[f]$ do not lie on a common cycle in $G'_{F_1,0}$. If either $[e]$ or $[f]$, say $[e]$, lies on a cycle in $G'_{F_1,0}$, then we proceed as above, but we need to choose $p_{\tilde{v}}$ so that the edges $[e_1], [e'_1]$ and $[e_3]$ have framework colour $[F_1]$, and $[e_2]$ has framework colour $[F_2]$ to guarantee that $G_{F_1,0}$ is a spanning unbalanced map graph of (G_0, ψ). If neither $[e]$ nor $[f]$ lie on a cycle in $G'_{F_1,0}$, then we need to distinguish two cases. Let C be the cycle in the connected component of $G'_{F_1,0}$ containing the edges $[e], [f]$. If there exists a path in $G'_{F_1,0}$ from a vertex in C to $[v_2]$ or $[v_3]$ that does not include $[v_1]$, then we may again proceed as above. Otherwise we choose a position for $p_{\tilde{v}}$ so that it is distinct from $\{p_w : w \in G'\}$, not fixed by $\tau(-1)$, and such that (G, p, θ, τ) is well-positioned and in (G_0, ψ) the edges $[e_1]$ and $[e'_1]$ have respective framework colours $[F_1]$ and $[F_2]$, and $[e_2]$ and $[e_3]$ both have framework colour $[F_1]$.

Suppose next that (G_0, ψ) is obtained from (G'_0, ψ') by a H3c move where $[v] \in G_0 \setminus G'_0$, and the H3c move deletes the unbalanced loop $[e] = [v_1][v_1]$ and the edge $[f] = [v_2][v_3]$ and adds the edges $[e_1] = [v][v_1]$ and $[e'_1] = [v][v_1]$ with respective gains $\gamma_1 = 1 \neq -1 = \gamma'_1$, and the edges $[e_2] = [v][v_2]$ and $[e_3] = [v][v_3]$ with respective gains γ_2 and γ_3. This case is completely analogous to the H3a case. If $[e]$ and $[f]$ belong to different induced monochrome subgraphs of G'_0, say $[e] \in G'_{F_1,0}$ and $[f] \in G'_{F_2,0}$, then we may choose $p_{\tilde{v}}$ so that $[e_1]$ and $[e'_1]$ have induced framework colour $[F_1]$, and $[e_2]$ and $[e_3]$ have induced framework colour $[F_2]$. If $[e]$ and $[f]$ belong to the same induced monochrome subgraph of G'_0, say $[e], [f] \in G'_{F_1,0}$, then Claim 4.7 applies, and we may choose $p_{\tilde{v}}$ so that $[e_1]$ and $[e'_1]$ have induced framework colour $[F_1]$, and $[e_2]$ and $[e_3]$ have respective framework colours $[F_1]$ and $[F_2]$, so that $G_{F_1,0}$ and $G_{F_2,0}$ are spanning unbalanced map graphs of (G_0, ψ).

If (G_0, ψ) is obtained from (G'_0, ψ') by a H3d move, then we may again proceed analogously to the H3a (or H3c) case. Note that if the loops $[e]$ and $[f]$ that are deleted in the H3d move are both in the same induced monochrome subgraph of G'_0, say $[e], [f] \in G'_{F_1,0}$, then they must lie in separate connected components of $G'_{F_1,0}$ (since they are both unbalanced cycles). So their removal results in two disjoint trees. With the addition of the vertex $[v]$ and the edges $[e_1], [e'_1], [e_2]$, these two trees are connected and a single (unbalanced) cycle is created in this connected component of $G_{F_1,0}$. □

Lemma 4.8. *Let (G_0, ψ) and (G'_0, ψ') be the gain graphs of the $\mathbb{Z}_2$-symmetric graphs (G, θ) and (G', θ'), respectively and suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a vertex-to-K_4 or vertex splitting move. If for (G', θ') there exists a realisation $\mathcal{G} = (G', p', \theta', \tau)$ which is well-positioned, C_2-symmetric and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$, then the same is true for (G, θ).*

Proof. By Theorem 4.2(A), there exists a well-positioned C_2-symmetric realisation p' of (G', θ') in

$(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$ so that the induced monochrome subgraphs $G'_{F_1,0}$ and $G'_{F_2,0}$ of (G'_0, ψ') are both spanning unbalanced map graphs. By Theorem 4.2(A), it now suffices to show that the vertex (or vertices) of $G_0 \setminus G'_0$ can be placed in such a way that the corresponding framework (G, p, θ, τ) is C_2-symmetric and well-positioned, and the induced monochrome subgraphs $G_{F_1,0}$ and $G_{F_2,0}$ are both spanning unbalanced map graphs in (G_0, ψ).

We fix two points x_1 and x_2 in the relative interiors of F_1 and F_2 respectively.

First we suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a vertex-to-K_4-move, where the vertex $[v]$ of (G'_0, ψ') (which may be incident to an unbalanced loop $[e]$) is replaced by a copy of K_4 with a trivial gain labelling (and $[e]$ is replaced by the edge $[f]$ with gain -1). Suppose without loss of generality that the loop $[e]$ (if present) has framework colour $[F_2]$. As shown in Figure 1, K_4 has a well-positioned placement in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$ where the two monochrome subgraphs are both trees. Moreover, we may scale this realisation so that all of the vertices of the K_4 lie in a ball of arbitrarily small radius. Let $B(p_{\tilde{v}}, r)$ be the open ball with centre $p_{\tilde{v}}$ and radius $r > 0$. Choose a placement of the representative vertices of the new K_4 to lie within $B(p_{\tilde{v}}, r)$ such that the vertices are distinct from $\{p_w : w \in V(G')\backslash\{\tilde{v}\}\}$, none of the vertex placements are fixed by $\tau(-1)$ and the resulting placement of the new K_4 is such that the monochrome subgraphs are both trees. If r is sufficiently small then the edge $[f]$ (if present) has the induced framework colour $[F_2]$ and all original edges retain their original colour. It can be assumed that the corresponding C_2-symmetric placement of G is well-positioned. Moreover, the induced monochrome subgraphs $G_{F_1,0}$ and $G_{F_2,0}$ of G_0 are clearly spanning unbalanced map graphs of (G_0, ψ).

Finally, we suppose that (G_0, ψ) is obtained from (G'_0, ψ') by a vertex split, where the vertex $[v]$ of (G'_0, ψ') (which is replaced by the vertices $[v_0]$ and $[v_1]$) is incident to the edge $[v][u]$ with trivial gain and the edges $[v][u_i]$, $i = 1, \ldots, t$, in G'_0. Without loss of generality we may assume that $[v][u] \in G'_{F_1,0}$. If we choose $p_{\tilde{v}_0} = p_{\tilde{v}}$ and $p_{\tilde{v}_1}$ to be a point on the line $L = \{p_{\tilde{v}} + tx_2 : t \in \mathbb{R}\}$ which is sufficiently close to $p_{\tilde{v}}$, then the induced framework colour for $[v_0][v_1]$ is $[F_2]$ and the induced framework colour for $[v_0][u]$ and $[v_1][u]$ is $[F_1]$. (Again we may assume the framework is well-positioned). Moreover, all other edges of (G'_0, ψ') which have been replaced by new edges in (G_0, ψ) clearly retain their induced framework colouring if $p_{\tilde{v}_1}$ is chosen sufficiently close to $p_{\tilde{v}}$. It is now easy to see that for such a placement of $\tilde{v}_0$ and $\tilde{v}_1$, both $G_{F_1,0}$ and $G_{F_2,0}$ are spanning unbalanced map graphs of (G_0, ψ). □

We are now ready to prove Theorem 4.3.

Proof. As mentioned earlier, (i) ⇒ (ii) follows from Corollary 2.20, and (ii) ⇒ (iii) follows from Theorem 3.20.

(iii) ⇒ (i): We employ induction on the number of vertices of G_0. By Theorem 4.2(A), for each of the base gain graphs there exists a representation $\tau : \mathbb{Z}_2 \to \mathrm{Isom}(\mathbb{R}^2, \|\cdot\|_\infty)$ and a realisation p such that $\mathcal{G} = (G, p, \theta, \tau)$ is well-positioned, C_2-symmetric and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_\infty)$, as indicated in Figure 1. (The two induced spanning map graphs $G_{F_1,0}$ and $G_{F_2,0}$ are shown in gray and black colour, respectively.) Since $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$ is isometrically isomorphic to $(\mathbb{R}^2, \|\cdot\|_\infty)$, there also exists a well-positioned, C_2-symmetric and χ_0-symmetrically isostatic realisation for each of the base graphs in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$.

Let $n \geq 5$ and suppose (i) holds for all gain graphs satisfying (iii) with at most $n - 1$ vertices. Let (G_0, ψ) have n vertices, and let (G'_0, ψ') be the penultimate graph in the construction sequence

of (G_0, ψ). By the induction hypothesis, there exists a realisation p' of the covering graph G' of (G'_0, ψ') in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$ so that (G', p', θ', τ) is well-positioned, C_2-symmetric and χ_0-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$.

If (G_0, ψ) is obtained from (G'_0, ψ') by a H1a, H1b, or H1c move, then the result follows from Lemma 4.4. If (G_0, ψ) is obtained from (G'_0, ψ') by a H2a, H2b, H2c, H2d or H2e move, then the result follows from Lemma 4.5. If (G_0, ψ) is obtained from (G'_0, ψ') by a H3a, H3b, H3c, or H3d move, then the result follows from Lemma 4.6. Finally, if (G_0, ψ) is obtained from (G'_0, ψ') by a vertex-to-K_4 or vertex splitting move, then the result follows from Lemma 4.8.

□

Next we establish the χ_1-symmetric counterpart to Theorem 4.3. The proof of this result is much simpler than the proof of Theorem 4.3 since the characterisation of $(2, 2, 2)$-gain-tight gain graphs in terms of a recursive construction sequence is significantly less complex than the one for $(2, 2, 0)$-gain-tight gain graphs.

Theorem 4.9. *Let $\|\cdot\|_{\mathcal{P}}$ be a norm on $\mathbb{R}^2$ for which $\mathcal{P}$ is a quadrilateral, and let (G, θ) be a $\mathbb{Z}_2$-symmetric graph. Further, let (G_0, ψ) be the gain graph for (G, θ). The following are equivalent.*

(i) There exists a C_2-symmetric realisation $\mathcal{G} = (G, p, \theta, \tau)$ which is well-positioned and χ_1-symmetrically isostatic in $(\mathbb{R}^2, \|\cdot\|_{\mathcal{P}})$;

(ii) (G_0, ψ) is $(2, 2, 2)$-gain tight;

(iii) (G_0, ψ) can be constructed from K_1 by a sequence of H1a,b moves, H2a,b moves, vertex-to-K_4 moves, and vertex splitting moves.

Proof. (i) ⇒ (ii): This follows again from Corollary 2.20.

(ii) ⇒ (iii): The proof proceeds by induction. Since (G_0, ψ) is $(2, 2, 2)$-gain-tight, it has no loops. If G_0 has a vertex $[v]$ of degree 2, then it is clearly admissible (via an inverse H1a or H1b move). So suppose G_0 has no degree 2 vertices. By $(2, 2, 2)$-gain-tightness, G_0 has a vertex $[v]$ of degree 3 with at least two neighbours. If $[v]$ has exactly two neighbours, then it is admissible (via an inverse H2b move, see also Lemma 3.10). Thus we may assume that every degree 3 vertex $[v]$ of G_0 has exactly three neighbours. It is easy to see that $[v]$ is admissible (via an inverse H2a move) unless it is contained in a balanced copy of K_4. (See also Lemma 3.9). The vertices of this K_4 cannot induce any additional edges since (G_0, ψ) is $(2, 2, 2)$-gain-tight. Denote this copy of K_4 as K. We may apply an inverse vertex-to-K_4 move unless there is a vertex $[x] \notin K$ and edges $[x][a]$ and $[x][b]$ with equal gains, where $[a], [b] \in K$. By switching, we may assume that both gains are 1. We may now apply an inverse vertex splitting move, contracting either $[x][a]$ and $[x][b]$, unless there exist vertices $[y]$ and $[z]$ that are distinct from the vertices of K and $[x]$ so that $[y][x]$ and $[y][a]$ are edges in (G_0, ψ) with the same gain, and $[z][x]$ and $[z][b]$ are edges in (G_0, ψ) with the same gain. By switching, we may again assume that the gains of these edges are all 1. We continue in this fashion, thereby constructing an increasing chain of subgraphs of (G_0, ψ) which are all $(2, 2, 2)$-gain tight and whose edges have all gain 1. (Note that at each step a new vertex is introduced for otherwise $(2, 2, 2)$-gain-sparsity is violated.) This sequence terminates after finitely many steps at which point there will be an admissible inverse vertex splitting move.

(iii) $\Rightarrow$ (i): Using Theorem 4.2(B), this result may be proved completely analogously to Theorem 4.3 (iii) $\Rightarrow$ (i). $\square$

5. Concluding remarks

One may be tempted to try to combine Theorems 4.3 and 4.9 to combinatorially characterise infinitesimal rigidity for half-turn symmetric frameworks. However this seems to be non-trivial. In particular, given a gain graph which contains a spanning $(2,2,2)$-gain-tight subgraph and a spanning $(2,2,0)$-gain-tight subgraph it is not clear that a placement exists that preserves both the colourings needed to apply Theorem 4.2(C).

It is also natural to try to extend Theorems 4.3 and 4.9 to higher-order groups, such as the cyclic group C_4 generated by a 4-fold rotation in the ℓ^1- or ℓ^∞-plane. In this case, Corollary 2.20 provides necessary gain-sparsity conditions for χ-symmetric infinitesimal rigidity. However, we are currently lacking analogues of Theorems 4.2(A) and 4.2(B) to prove the sufficiency of these counts.

There is a second form of vertex splitting, known as the vertex-to-4-cycle move [10, 14], which one may use instead of vertex splitting to give analogous inductive constructions to Theorem 3.20 and Theorem 4.9 (ii) $\Leftrightarrow$ (iii). In fact, in the case of $(2,2,0)$-gain-tight gain graphs, this alternative gives a non-trivial simplification to the proof, replacing the maximal balanced triangle sequence considerations with a direct counting argument. However in both the symmetric and anti-symmetric contexts this construction operation does not seem to be amenable to finding appropriate rigid placements.

In [15] symmetric rigidity is considered for frameworks in Euclidean space that are restricted to move on a fixed surface. In particular the matroidal classes of $(2,2,2)$, $(2,2,1)$ and $(2,2,0)$-gain-tight gain graphs are the relevant sparsity types for frameworks restricted to an infinite circular cylinder. Hence our recursive construction of $(2,2,0)$-gain-tight gain graphs may be useful in establishing an analogue of Theorem 4.3 for the appropriate symmetry group, that is for half-turn symmetric frameworks on the cylinder with rotation axis perpendicular to the axis of the cylinder.

[1] C. D'Ambrosio and L. Liberti, Distance geometry in linearizable norms. *Geometric science of information*, 830–837, Lecture Notes in Comput. Sci., 10589, *Springer, Cham,* 2017.

[2] S. Fiorini, T. Huynh, G. Joret, A. Varvitsiotis, The excluded minors for isometric realizability in the plane. SIAM J. Discrete Math. 31 (2017), no. 1, 438–453.

[3] R. Ikeshita and S. Tanigawa, *Count matroids of group-labeled graphs*, Combinatorica **38**, (2018), 1101–1127.

[4] T. Jordán, V. Kaszanitzky and S. Tanigawa *Gain-sparsity and symmetry-forced rigidity in the plane*, Discrete Comput. Geom. **55**, (2016), 314–372.

[5] D. Kitson and S.C. Power, *Infinitesimal rigidity for non-Euclidean bar-joint frameworks*, Bull. London Math. Soc. **46** (4), (2014), 685–697.

[6] D. Kitson and B. Schulze, *Maxwell-Laman counts for bar-joint frameworks in normed spaces*, Linear Algebra and its Applications **481**, (2015), 313–329.

[7] D. Kitson and B. Schulze, *Motions of grid-like reflection frameworks*, Journal of Symbolic Computation **88**, (2018), 47–66.

[8] G. Laman, *On graphs and rigidity of plane skeletal structures*, J. Engrg. Math. **4**, (1970), 331–340

[9] A. Lee and I. Streinu, *Pebble game algorithms and sparse graphs*, Discrete Math., **308** (8), (2008), 1425–1437.

[10] L. Lomeli, L. Moshe and W. Whiteley, *Bases and circuits for 2-rigidity: constructions via tree partitions*, Technical report, http://www.math.yorku.ca/Who/Faculty/Whiteley/menu.html.

[11] J. Malestein and L. Theran, *Frameworks with Forced Symmetry I: Reflections and Rotations*, Discrete Comput. Geom. **54** (2), (2015), 339–367.
[12] A. Nixon and J. Owen, *An inductive construction of* $(2,1)$*-tight graphs*, Contributions to Discrete Mathematics, **9** (2), (2014), 1–16.
[13] A. Nixon, J. Owen and S.C. Power, *Rigidity of frameworks supported on surfaces*, SIAM Journal on Discrete Mathematics, **26** (4), (2012), 1733–1757.
[14] A. Nixon, J. Owen and S.C. Power, *A characterization of generically rigid frameworks on surfaces of revolution*, SIAM Journal on Discrete Mathematics, **28** (4), (2014), 2008–2028.
[15] A. Nixon and B. Schulze, *Symmetry-forced rigidity of frameworks on surfaces*, Geometriae Dedicata **182** (1), (2016), 163–201.
[16] B. Schulze, *Block-diagonalized rigidity matrices of symmetric frameworks and applications*, Contributions to Algebra and Geometry **51** (2), (2010), 427–466.
[17] B. Schulze and S. Tanigawa, *Infinitesimal rigidity of symmetric bar-joint frameworks*, SIAM Journal on Discrete Mathematics **29** (3), (2015), 1259–1286.
[18] B. Schulze and W. Whiteley, *Rigidity and scene analysis*, in Handbook of Discrete and Computational Geometry, C.D. Toth, J. O'Rourke, J.E. Goodman, editors, Third Edition, Chapman & Hall CRC, 1593–1632, 2017.
[19] M. Sitharam, A. St. John, J. Sidman, Handbook of Geometric Constraint Systems Principles, Chapman & Hall CRC, 2018.
[20] W. Whiteley, *Some Matroids from Discrete Applied Geometry*, Contemporary Mathematics, AMS **197** (1996), 171–311.
[21] T. Zaslavsky, *Biased graphs. I. Bias, balance, and gains*, J. Comb. Theory Ser. B **47** (1), (1989), 32–52.