BIASED GRAPHS. IV. GEOMETRICAL REALIZATIONS

THOMAS ZASLAVSKY

Abstract. A *gain graph* is a graph whose oriented edges are labelled invertibly from a group $\mathfrak{G}$, the *gain group*. A gain graph determines a biased graph and therefore has three natural matroids (as shown in Parts I–II): the *bias matroid* G has connected circuits; the *complete lift matroid* L_0 and its restriction to the edge set, the *lift matroid* L, have circuits not necessarily connected. We investigate representations of these matroids. Each has a canonical vector representation over any skew field F such that $\mathfrak{G} \subseteq F^*$ (in the case of G) or $\mathfrak{G} \subseteq F^+$ (in the case of L and L_0). The representation of G is unique up to change of gains when the gain graph is full, but not in general. The representation of G or L is unique or semi-unique (up to changing the gains) for 'thick' biased graphs. The lift representations are unique (up to change of gains) for L_0 but not for L. The bias matroid is representable also in other ways by points and hyperplanes; one of these representations dualizes the vector representation, while two, in projective space, strongly generalize the theorems of Menelaus and Ceva. (The latter specialize to properties of the geometry of midpoints and farpoints, and of median and edge-parallel hyperplanes, in an affine simplex.) The dual hyperplane representation can be abstracted away from fields to a kind of equational logic and permutation geometry that exist for every gain group. The lift matroids are representable by orthographic points and by linear, projective, and affinographic hyperplanes. L_0 also has a metric hyperplanar representation that depends on the Pythagorean theorem. Incidental results are that Whitney's 2-isomorphism operations preserve gains and, to an extent, matroids; and new definitions of the bias and lift matroids based on extremal properties of the rank functions.

Date: Version of August 18, 2002, as trivially revised August 8, 2003.

2000 *Mathematics Subject Classification. Primary* 05B35, 05C22; *Secondary* 03B99, 05A15, 05C50, 51A45, 52C35.

Key words and phrases. Biased graph, gain graph, balance, bias matroid, graphic lift matroid, matroid representation, arrangement of hyperplanes, hyperplane separation, Menelaus, Ceva, orthographic representation, Pythagorean representation, chromatic polynomial, matroid characteristic polynomial, Poincaré polynomial, Whitney-number polynomial, logic of equations, permutation representation, permutation gain graph, Whitney 2-isomorphism, thick biased graph, geometric semilattice.

Research substantially assisted by grants from the National Science Foundation: in 1976–1977 at the Massachusetts Institute of Technology, in 1984–1985 while I was visiting the University of Evansville, and in 1986 and in 2000–2002 (grant DMS-0070729) at Binghamton.

Acknowledgements. I thank the referee of the original version for pointing out the need for numerous examples, of which there had been none in this series. I thank Marge Pratt for enormous typing assistance over innumerable years.

Contents

Introduction

This article concerns a meeting point of graph theory, matroid theory, and geometry: the geometrical and equational representation theory of matroids of gain graphs. A *gain graph* Φ is a graph together with a mapping (the *gain function*) from the edge set E to a group (the *gain group*), such that reversing the direction of an edge inverts the gain. We call a circle (edge set of a simple closed path) *balanced* if its edge gains, taken in circular order, multiply to the identity. This gives a *biased graph*: a graph together with a subclass of its circles, called the 'balanced' circles, such that if the union of two balanced circles is a theta graph, the third circle in their union is also balanced. The matroids we investigate are three matroids associated with any biased graph, and in particular with any gain graph Φ. Say a subgraph is *balanced* if every circle in it is balanced (and it contains no half edges—see Section 1). For an edge set S let $b(S)$ denote the number of balanced components and $c(S)$ the total number of components of (N, S), where N is the node set of Φ; and let $n = \#N$. The *bias rank* of S is

$$\mathrm{rk}_G\, S = n - b(S)$$

(provided n is finite; in general see Section 1); it is the rank function of the *bias matroid* $G(\Phi)$, whose point set is E. Let e_0 be an *extra* or *ideal point* not in E and define the *lift rank*

$$\mathrm{rk}_L\, S = \begin{cases} n - c(S) & \text{if } S \subseteq E \text{ is balanced,} \\ n + 1 - c(S) & \text{if } S \text{ is unbalanced or } e_0 \in S \end{cases}$$

(again, provided n is finite). This is the rank function of the *complete lift matroid* $L_0(\Phi)$, whose ground set is $E \cup \{e_0\}$; its restriction to E is the *lift matroid*. Our fundamental results (in Sections 2 and 4) are the construction of 'canonical' representations of each of these matroids as linear dependence matroids of vectors over a skew field which contains the gain group as (in the bias case) a multiplicative subgroup or (in the lift case) an additive subgroup. For any gain group we also give a vector-like permutation representation of $G(\Phi)$ and a closely related but purely logical representation as a system of two-term equations (Section 3).

We also demonstrate various refinements of the basic vector representations. For instance, dualizing gives a representation by hyperplanes whose equations are (in the bias case) of the form $x_i = \alpha x_j$ or $x_i = 0$ or (in the lift case) of the form $x_j - x_i = \alpha x_0$ or $x_0 = 0$; we give real and complex examples of hyperplanar bias representations, with the number of regions in the real case and the Betti numbers of the complement in the latter. (A further duality, which we omit, would give zonotopal representations and related enumerations.) Projecting the vector representation of $G(\Phi)$ yields an extension of the theorem of Menelaus, more general than the higher-dimensional Menelaus theorems of [7, 10, 23]. Dually, we have a very general Ceva theorem. In these results we take an affine basis B and arbitrary points, called *apices*, on the lines generated by B in projective space. A *Cevian* or *apical hyperplane* is the span of an apex and the basis elements not on its line. The Menelaus theorem describes the projective dependencies of the apices. The Ceva theorem describes the intersection pattern of the apical hyperplanes. Our theorems seem to be ultimate of this type because the apices are subject to no restriction. (However, there do appear to be other generalizations, at least of Ceva: see Boldescu [7]. It would be interesting to see a gain-graphic formulation and a Menelæan analog of Boldescu's theorem.) In order to understand these representations and several examples involving midpoints and medians of edges, as well as infinite points on and parallel hyperplanes to edge lines, we are forced to develop some tools of projective geometry that appear to be obscure or maybe new. For the lift and complete lift matroids there are an affine variant of canonical representation of $L(\Phi)$, which we call an 'orthographic' representation, and a 'Pythagorean' representation of $L_0(\Phi)$ by Euclidean hyperplanes. In a subsequent paper (Part VI) we give a different presentation of some of these representations in the style of synthetic geometry.

We consider the possibility of uniqueness theorems, asserting that any linear representation of a gain-graphic or biased-graphic matroid is of canonical type. That is not true in general. We prove it true for natural large, though specialized, classes of gain and biased graphs (see especially Theorem 7.1); but we do not solve the important problem of determining the gain or biased graphs having only canonical representations.

Section 5 on Whitney operations and separable graphs casts light on some simple aspects of geometric representation. Remarkably, although the bias matroid is not invariant under Whitney's 2-isomorphism operations, the existence of a canonical representation over a given skew field is. Section 6 on alternative definitions of the lift and bias matroids, required here for the proof of Theorem 7.1, is most interesting as a continuation of Part II. Section 8 concludes the treatment of the biased K_4's in Parts I, II, and III. Section 9 suggest some final research questions.

The main questions raised by this work are based on our concept of canonical representation. We ask: which biased graphs have noncanonical representations, and in which skew fields does a biased graph Ω have additive or multiplicative gains. We collect some remarks about these properties in Section 9. I hope these and other gaps in our understanding of biased graphs will be taken as a challenge to discover answers—and, no doubt, further questions.

The numerous examples, some illustrated, include the spikes, swirls, and whirls of matroid representation theory, Menelæan and Cevian corollaries about the affine geometry of a simplex, and many others intended to illuminate the main results. In subsequent parts [48] we fill out the theory with a much wider variety of examples. Dowling geometries in their gain-graphical aspect, and various generalizations, have interesting features related to unique and semi-unique representability and isomorphism (Part V). We have a good deal to say about contrabalanced graphs (Part VII), where no circle is balanced; to mention only one of many questions, we are led to speculate about a contrabalanced analog of Tutte's theory of nowhere-zero flows. Part VII is a grab-bag of investigations of other special types of bias. About balanced graphs, simple as they are, there is a little to say. On additively biased graphs—equivalently, signed graphs—we review some representation results that are known, but not all well known, looking in particular at antibalanced biased graphs (equivalently, all-negative signed graphs). We study two kinds of bias based on orientation of the graph: bias from poise, where a balanced circle has no majority direction, and antidirection bias, where directions reverse at every step around a balanced circle. Bias from Hamiltonian circles gives us an excellent opportunity to explore the relationship between gain groups and 'canonical' matroid representations, although we can say little about noncanonical representations.

If the gain graph is finite and the scalar field is an ordered one, say the real numbers, we can define hyperplane separations of a point representation and faces (k-dimensional cells for any k) of a hyperplane representation. Basic geometric theorems and the chromatic theory of Part III enable us to express the number of hyperplane separations and the numbers of

k-dimensional faces in terms of chromatic invariants of the gain graph. In the infinite case one cannot count, but one can still ask the basic question: what is the abstract structure of the hyperplane separations of a vector representation, or dually of the faces of a hyperplanar representation, for a gain graph with ordered gain group, and more generally for a biased graph? That means finding the correct definition of orientation of a biased graph so as to be compatible with the oriented matroid [6, 5] of the canonical vector and hyperplanar representations of a real multiplicative or additive gain graph—that is, of $G(\Phi)$ when $\mathfrak{G} \subseteq \mathbb{R}^*$ or $L(\Phi)$ when $\mathfrak{G} \subseteq \mathbb{R}^+$ (Theorems 2.1 and 4.1 and Corollary 2.2). The definition should be compatible with that of signed-graph orientation in [43]. In the first version of this article, in 1986, I could only speculate on a possible approach. Recently Slilaty solved and generalized the first half of this problem, finding the correct definition of oriented-matroid cycles of the bias matroid of a biased graph with signed edges; that and more will be found in his doctoral dissertation [31].

This article is a continuation of the theory of biased graphs in Parts I–III [41]. Nevertheless it should be possible to read it without being familiar with the previous parts. For still more work on gain and biased graphs the reader may peruse [45].

Since this part is long, some guidance to the reader may be appropriate. The core is Sections 2.1 and 4.1, followed by Section 7. After these in importance are Sections 2.5, 2.6, and 4.5, and then the remainder of the paper.

1. Preliminaries

In principle we assume that the reader is acquainted with relevant definitions from Parts I–III, but in practice that will not usually be necessary. Here we re-emphasize some standard notation and introduce some new notation and concepts. We cite earlier and later parts [41, 48] in the style 'Section I.2', 'Theorem III.5.1'.

Always, Γ is a graph (N, E); $n = \#N$, which may be infinite; Ω is a biased graph $(\Gamma, \mathcal{B})$ (Section I.1); $\mathfrak{G}$ is a group; and Φ is a gain graph $(\Gamma, \varphi, \mathfrak{G})$, meaning that it has underlying graph Γ, gain group $\mathfrak{G}$, and gain function φ (Section I.5). We may leave the group implicit, writing $\Phi = (\Gamma, \varphi)$. We sometimes write $\|\Omega\|$ or $\|\Phi\|$ for the underlying graph Γ. E_* is the set of *ordinary edges* (links and loops). Other edges are *half edges* (one endpoint) and *loose edges* (no endpoints, but we shall have little use for them here). E_0 is $E \cup \{e_0\}$, where the *extra point* $e_0 \notin E$. For the gain of an edge e with endpoints v and w (written $\nu_\Gamma(e) = \{v, w\}$, a multiset since v and w may be equal) we usually write $\varphi(e; v, w)$ in order to indicate the sense in which the gain is measured; thus $\varphi(e; w, v) = \varphi(e; v, w)^{-1}$. The biased graph determined by Φ is denoted by $\langle\Phi\rangle$ (Section I.5; the former notation $[\Phi]$ was an error and should be reserved for switching classes). If Ω has the form $\langle\Phi\rangle$ and $\mathfrak{G}$ is the gain group of Φ, we say that Ω *has gains in* $\mathfrak{G}$. The bias, lift, or complete lift matroid of Φ is that of $\langle\Phi\rangle$; we write $G(\Phi)$ for $G(\langle\Phi\rangle)$, etc. These matroids are defined by the rank functions stated in the introduction when the biased graph has finite order; in general we define the rank of S in terms of the edge-induced subgraph $\Omega{:}S$, which is the biased subgraph consisting of S and its incident node set $N(S)$; then the rank functions are:

$$\operatorname{rk}_G S = \#N(S) - b(\Omega{:}S)$$

for the bias matroid and

$$\operatorname{rk}_L S = \begin{cases} \#N(S) - c(\Omega{:}S) & \text{if } S \subseteq E \text{ is balanced,} \\ \#N(S) + 1 - c(\Omega{:}S) & \text{if } S \text{ is unbalanced or } e_0 \in S \end{cases}$$

for the lift and complete lift matroids, with the proviso that balanced loops are to be ignored in calculating $N(S)$ and $\Omega{:}S$. (These definitions coincide with those in the introduction when n is finite. We interpret both ranks as ∞ when $N(S)$ is infinite.) An *unbalanced figure* is an unbalanced circle or a half edge.

The *restriction* $\Phi|(W,S)$ or $\Omega|(W,S)$ of a gain or biased graph to a subgraph (W,S) is the subgraph with gains restricted to S or with balanced circles equal to those contained in S. Restriction to S, as in $\Omega|S$, means restriction to (N,S) (as opposed to an *edge-induced* subgraph $\Omega{:}S$, which is the restriction to $(N(S),S)$). The edge set *induced* by a node set W is $E{:}W = \{e \in E : \varnothing \neq N(e) \subseteq W\}$; the subgraph induced by W is $\Omega{:}W = \Omega|(W, E{:}W)$ and similarly for graphs and gain graphs.

A *switching function* η is any function $N \to \mathfrak{G}$. The *switching of* Φ *by* η is Φ^η whose gain function is $\varphi^\eta(e;v,w) = \eta(v)^{-1}\varphi(e;v,w)\eta(w)$. Two gain graphs are *switching equivalent* if one is a switching of the other. The switching equivalence class, or *switching class*, of Φ is denoted by $[\Phi]$. Since the balanced circle class of a gain graph is unaltered by switching, the various matroids of the gain graph are also unchanged.

A *potential* for a balanced gain graph Φ is a function $\theta : N \to \mathfrak{G}$ such that $\varphi(e;v,w) = \theta(v)^{-1}\theta(w)$ for each edge $e{:}vw$. (A potential cannot exist if Φ has an unbalanced circle.) The reciprocal θ^{-1}, defined by $\theta^{-1}(v) = (\theta(v))^{-1}$, is a switching function that switches Φ to the identity gain graph $(\|\Phi\|, 1)$. Conversely, any such switching function is the reciprocal of a potential. An example of a potential in a connected gain graph is obtained by choosing a root node v and defining $\theta(w) = \varphi(P_{vw})$ if P_{vw} is a vw-path.

If X is a subset of a linear, affine, or projective space Λ over some skew field, we write $\operatorname{span} X$ for the subspace (= flat) spanned by X. We write $M(X)$ for the matroid of X under the appropriate dependence relation. (Note that we distinguish between Λ and its matroid $M(\Lambda)$.) If $\mathcal{H}$ is a family of hyperplanes in Λ, a *flat of* $\mathcal{H}$ is a subspace formed by intersecting members of $\mathcal{H}$, for example the whole space, which is $\bigcap \varnothing$; but we exclude the empty subspace in the affine case. In a projective space an *affine flat* is a flat not contained in the ideal hyperplane h_∞.

Let M be a matroid on ground set E. M is a *line* if it has rank 2. $\operatorname{Lat} M$ is the lattice of flats. $\operatorname{Lat}^{\mathrm{b}} \Phi$ is the semilattice of balanced flats in $G(\Phi)$ or $L(\Phi)$.

Now, let F be a skew field. A *linear* (i.e., *vector*) *representation* of M *over* F is a mapping f of E into the point set of a linear space Λ over F (with F, if not commutative, acting on the left unless otherwise stated) such that, for each $S \subseteq E$, $f(S)$ is dependent exactly when S is. An *affine* [or, *projective*] *representation* of M is a mapping f of $E \setminus \{\text{loops}\}$ into the point set of an affine [or, projective] space Λ such that, for each $S \subseteq E$, $f(S)$ is dependent exactly when S is. Λ need not be coordinatizable, but if it has coordinates in a skew field F we say f is an affine [or, projective] representation of M over F. Similarly, a *hyperplane representation* of M is a mapping from E to a family of hyperplanes in Λ which preserves dependence and independence. (For an affine hyperplane representation this definition is not quite correct, but it is correct to think of it as a projective hyperplane representation, not using the ideal hyperplane, which is then restricted to affine space.) Our definition of projective space is broad: we admit a line (of order at least 2); we consider it desarguesian when it is coordinatized by a skew field.

Two representations of a matroid M in a linear or projective space Λ are *projectively equivalent* if they are related by a projective automorphism of Λ. In a linear space or a coordinatized projective space a projective automorphism is a linear operator combined with a field automorphism and scaling of vectors (multiplication by a nonzero scalar). In a projective space of dimension 2 or higher, a projective automorphism is a collineation. (We do not define projective equivalence in a noncoordinatized projective line.) We say M has *(projectively) unique representation* in Λ if it has only one representation in Λ up to projective equivalence. All graphic matroids have projectively unique representation [9].

Suppose E is a nonempty set of points in a real linear or affine space and M is its linear dependence matroid. A *hyperplane separation* of E is a partition of E into at most two parts that are separable by a hyperplane. The number of hyperplane separations of E is $\frac{1}{2}|p_M(-1)|$, where $p_M(\lambda)$ is the characteristic polynomial of M (Section III.5). This is known by [36, Corollaries 6.1 and 6.2], also by [26, Theorem 3.1 and note on p. 243] (and see [25]).

An *arrangement of hyperplanes* is a finite family of hyperplanes in d-dimensional linear, affine, or projective space over some skew field. The characteristic polynomial of $\mathcal{H}$ is $p_{\mathcal{H}}(\lambda) = \sum_t \mu(\varnothing, t)\lambda^{\dim t}$, where t ranges over all intersections (nonvoid, except in the projective case) of members of $\mathcal{H}$ and μ is the Möbius function of the semilattice of intersections (see [36]). If $\mathcal{H}$ is linear and M is the matroid of $\mathcal{H}$, then $p_{\mathcal{H}}(\lambda) = \lambda^{d-\operatorname{rk} M} p_M(\lambda)$; if $\mathcal{H}$ is projective, then $p_{\mathcal{H}}(\lambda) = \lambda^{d+1-\operatorname{rk} M} p_M(\lambda)$; there is a similar formula for affine arrangements.

Suppose $\mathcal{H}$ is a real hyperplane arrangement. $\mathcal{H}$ decomposes the space into cells of various dimensions, called the *faces* of $\mathcal{H}$. Let f_k be the number of k-dimensional faces. The *face generating polynomial* is $f_{\mathcal{H}}(x) = \sum_k x^{d-k} f_k$. It is known from [36, Theorems A and B] that $f_{\mathcal{H}}(x) = (-1)^{\operatorname{rk} M} w_M(-x, -1)$ in the linear case, $\frac{1}{2}\left[(-1)^{\operatorname{rk} M} w_M(-x, -1) + x^{\operatorname{rk} M}\right]$ in the projective case, where $w_M(x, \lambda)$ is the Whitney-number polynomial of M (Section III.5). There are similar formulas in the real affine case both for all faces and for bounded faces [36, Theorems A and C].

If $\mathcal{H}$ is a complex hyperplane arrangement in the affine space $\mathbb{A}^d(\mathbb{C})$, its complement X has complicated topology, but the Poincaré polynomial of X, $P_{\mathcal{H}}(y) = \sum_{i\geq 0} \operatorname{rk} H^i(X, \mathbb{Z}) y^i$, is computable from $\mathcal{H}$: it equals $(-y)^d p_{\mathcal{H}}(-1/y)$ ([28], or see [29, Theorem 5.93]).

Thus when we can calculate the characteristic polynomial of a gain-graphic or biased-graphic matroid we can find the number of hyperplane separations, or of d-faces, of a real representation; if we can calculate the Whitney-number polynomial we can find the number of faces of each dimension of a hyperplane representation. Moreover, the individual coefficients of these polynomials have geometrical interpretations, both in the complex (as above) and real [18] cases. In Part III we discussed methods of evaluating the polynomials and applied them to several interesting types of examples. In this and future parts we apply the conclusions herein to study the geometry of many of those examples and interpret their numbers.

2. Geometry of the bias matroid

A *bias representation* of Ω is a vector, affine, or projective representation of $G(\Omega)$. Amongst all possible bias representations we single out one type that is canonical for a gain graph.

2.1. **Canonical representations.**

Suppose the gain group $\mathfrak{G}$ of Φ is a multiplicative subgroup of F. (We write $\mathfrak{G} \leq F^*$.) Let Λ be an F-vector space and $N \to \Lambda$ (written $v \mapsto \hat{v}$) a bijection of N onto an independent subset $\hat{N}$ of Λ. For each edge we define a vector $x_\Phi(e)$, or simply $x(e)$, $\in \Lambda$ by

$$x_\Phi(e) = \begin{cases} -\hat{v} + \varphi(e; v, w)\hat{w} & \text{if } \nu_\Gamma(e) = \{v, w\}, \\ \hat{v} & \text{if } e \text{ is a half edge at } v, \\ 0 & \text{if } e \text{ is a loose edge or balanced loop.} \end{cases}$$

This definition for a link or loop corresponds to orienting e from v to w; the opposite orientation would yield $x(e) = \varphi(e; w, v)\hat{v} - \hat{w}$. But the latter is a non-zero multiple of the former; consequently $x(e)$ is well defined up to non-zero scalar multiplication and it is completely well defined if every link and loop of Γ is given a direction. For notational convenience we define $x(S) = \{x(e) : e \in S\}$ (which may be a multiset) if $S \subseteq E$ and $\hat{W} = \{\hat{w} : w \in W\}$ for $W \subseteq N$. We call x_Φ a *standard bias representation of* Φ.

The reader may now go directly to part (a) of Theorem 2.1. The next paragraphs prepare for part (b).

Let $\mathbf{I}(\Phi)$ be the $N \times E$ matrix whose columns are the vectors $x(e) \in F^N$; we call this an *incidence matrix* of Φ (associated with the particular orientation used to define x; thus $\mathbf{I}(\Phi)$ is well defined up to nonzero scalar multiplication of columns).

We call a mapping $f : E \to \Lambda$ a *canonical bias representation of* Φ (in full, a *canonical linear*, or *vector*, *bias representation over* F) if there exists a mapping $N \to \Lambda$ of N onto an independent set $\hat{N}$ under which each $f(e)$ is a non-zero multiple of $x(e)$, or if Λ can be enlarged so that such a mapping exists. From the projective viewpoint (anticipating Theorem 2.1(a)), f is an embedding of $G(\Phi)$ in the projective space $\mathbb{P} = (\Lambda \setminus \{0\})/F^*$ that is completely determined once $\hat{N}$ is chosen. (We must remember that a balanced loop or loose edge, which in a vector representation corresponds to the zero vector, in a projective or affine representation is represented by no point at all.) If $\Lambda = \operatorname{span} \hat{N}$ we can regard Λ as the coordinatized vector space F^N and $\mathbb{P}$ as the particular projective space $\mathbb{P}^{N-1} = (F^N \setminus \{0\})/F^*$, with coordinates x_v for $v \in N$. (Although N is not a number, we write a superscript $N - 1$ to suggest the dimension and coordinate system of the space. If N is finite, the dimension equals $n - 1$.)

Suppose Λ and $\hat{N}$ as in the preceding paragraph. *Scaling* Λ by a switching function $\eta : N \to F^*$ means multiplying each vector $\hat{u} \in \hat{N}$ by the arbitrary nonzero scalar $\eta(u)$. This has the effect of transforming a vector $\xi = \sum_{u \in N} \xi_u \hat{u}$ to $\xi^\eta = \sum_{u \in N} \xi_u \eta(u)\hat{u}$; thus, $(\xi^\eta)_u = \xi_u \eta(u)$. Scaling is defined only for vectors spanned by $\hat{N}$. If $f : E \to \Lambda$ is any function whose image lies in $\operatorname{span} \hat{N}$, we define f^η by $f^\eta(e) = f(e)^\eta$; i.e., $f^\eta(e)_u = f(e)_u \eta(u)$ for each $u \in N$.

We call $f : E \to \Lambda$ a *canonical bias representation of the switching class* $[\Phi]$ if f is a canonical bias representation of a gain graph $\Phi^\eta \in [\Phi]$. In terms of linear algebra, f is virtually any function obtained from x_Φ through scaling Λ. Any canonical bias representation of $[\Phi]$ is identical to a scaling of a switching of a canonical bias representation of Φ. This is the meaning of the relationship

$$x_{\Phi^\eta}(e; v, w) = \eta(v)^{-1}(x_\Phi)^\eta(e; v, w). \tag{2.1}$$

Thus switching gains gives a new but projectively equivalent representation; so a switching class $[\Phi]$ has a unique canonical bias representation up to projective equivalence. (The converse is not true; see Section 2.3.)

We have so far defined a canonical bias representation of a gain graph and a switching class. A definition for biased graphs will appear later.

A function $f : E \to F^N$ and subset $S \subseteq E$ determine functions $f\big|_S$, the restriction of f to S, and in a complicated way $f/S : S^c \to F^{\pi_b(S)}$, defined on $S^c = E \setminus S$. To define f/S, let $S = S_0 \cup S_1 \cup \cdots \cup S_k$ where S_0 is the union of all unbalanced components of S and the other S_i are the balanced components of S. Let η be a switching function under which $\varphi^\eta\big|_{S_i} \equiv 1$ for each $i > 0$. For $e \in S^c$ and $V \in \pi_b(S)$, set $(f/S)(e)_V = \sum_{u \in V} f^\eta(e)_u$. This definition of f/S is unique up to switching of Φ/S so it is well defined on $[\Phi/S]$. In fact, f/S is the image of f under a vector space homomorphism. In $F^{\pi_b(S)}$ let $\hat{\hat{V}}$ be the unit vector in the V-direction, and let $\hat{\hat{\pi}}_b(S)$ be the basis composed of all such unit vectors. The homomorphism sends $\hat{u} \in \hat{N}$ to $\hat{\hat{V}} \in \hat{\hat{\pi}}_b(S)$ if $u \in V \in \pi_b(S)$ and to 0 if no such V exists.

Theorem 2.1. *Let $\Phi = (\Gamma, \varphi, \mathfrak{G})$ be a gain graph and F a skew field containing $\mathfrak{G}$ as a multiplicative subgroup.*

(a) *The linear dependence matroid of the vectors $x_\Phi(e)$ is isomorphic, under the mapping x_Φ, to the bias matroid $G(\Phi)$.*

(b) *For any canonical bias representation f and any $S \subseteq E$, $f\big|_S$ is a canonical bias representation of $\Phi|S$ and f/S is a canonical bias representation of $[\Phi/S]$.*

Proof of (a). We may treat a half edge as an unbalanced loop (with gain zero, say).

It is easy to show by trimming pendant edges that $x(S)$ is independent if S is a forest.

Suppose S is a circle, say $S = e_1 e_2 \cdots e_k$ with e_i oriented from v_{i-1} to v_i. (We take subscripts modulo k.) A linear relation $\sum_1^k \alpha_i x(e_i) = 0$ requires that $\alpha_{j+1} = \alpha_j \varphi(e_j)$ for each j. Therefore, $\alpha_1 \equiv \alpha_{k+1} = \alpha_k \varphi(e_k) = \cdots = \alpha_1 \varphi(e_1) \cdots \varphi(e_k) = \alpha_1 \varphi(S)$. We have $\alpha_1(1 - \varphi(S)) = 0$; so either S is balanced and its vectors are dependent, or it is unbalanced and they are independent.

In the latter case, letting $\alpha_1 = 1$ and $\alpha_j = \varphi(e_1) \cdots \varphi(e_{j-1})$ for $2 \leq j \leq k$, we have $\sum_1^k \alpha_i x(e_i) = (\varphi(S) - 1)\hat{v}_k$. Therefore $\hat{v}_k$ and consequently all $\hat{v}_i$, $i \leq k$, are in $\operatorname{span} x(S)$. That is, $\operatorname{span} x(S) = \operatorname{span} \hat{N}(T)$. If we enlarge S to a connected edge set T, we have $\operatorname{span} x(T) = \operatorname{span} \hat{N}(T)$.

One can easily deduce that, if R is an edge set each of whose components has cyclomatic number at most one and contains no balanced circle, then R is independent. But if a component of R contains a balanced circle or has cyclomatic number at least two, R is dependent. It follows that the linear dependence matroid of $x(E)$ is $G(\Phi)$. □

Proof of (b). The statement for $f\big|_S$ is trivial.

For f/S it suffices to consider the possible types of edge $e \in E(\Phi \setminus S)$.

First, e may have two endpoints (not necessarily different) in balanced components (not necessarily different) of Φ/S. Say the endpoints are v and w and $v \in V$, $w \in W$, with $V, W \in \pi_b(S)$. We have $f(e) = \alpha_e x_\Phi(e)$, whence $f^\eta(e) = \alpha_e \eta(v) x_{\Phi^\eta}(e) = \alpha_e \eta(v)[-\hat{v} + \varphi^\eta(e; v, w)\hat{w}]$, so that

$$(f/S)(e) = \alpha_e \eta(v)[-\hat{\hat{V}} + \varphi^\eta(e; v, w)\hat{\hat{W}}] = \alpha_e \eta(v) x_{\Phi^\eta/S}(e).$$

Second, e may have two endpoints of which only one, or neither, is in a balanced component of S. Thus e becomes a half or loose edge, respectively, in Φ/S. Or, e may be a half or loose edge in Φ. All these cases are easy to verify.

The conclusion is that f/S is a canonical bias representation of Φ^η/S. $\square$

In the case of a commutative field we can say more. Let C be a circle in Φ; let $\mathbf{I}(\Phi)$ be an incidence matrix of Φ in which the edges of C are all oriented the same way around C, and let M be the square submatrix of $\mathbf{I}(\Phi)$ indexed by $N(C)$ and C (that is, an incidence matrix of $\Phi{:}N(C)|C$). Then

$$\det M = \pm(1 - \varphi(C)). \tag{2.2}$$

We omit the proof.

Dowling [13, Theorem 10] gave a proof of Theorem 2.1(a) for $\Phi = \mathfrak{G}K_n^\bullet$ (the full $\mathfrak{G}$-expansion of K_n, defined in Examples I.6.7 and III.3.7 and Part V) along with a strong converse for that example ([13, Theorem 9], generalized in our Proposition 2.4). His result implies Theorem 2.1(a) for finite order, since the bias matroid of any gain graph of finite order is essentially a submatroid of $G(\mathfrak{G}K_n^\bullet)$. His proof is based on a presentation of $G(\mathfrak{G}K_n^\bullet)$ as a kind of partition lattice; our proof, based on a graphical presentation, is much simpler. Special cases have appeared as well. Proofs for $\mathfrak{G} = \mathbb{R}^*$ have appeared in the literature of networks with gains; cf. [27]. A detailed proof for $\#\mathfrak{G} = 2$ in [37, Section 8A] is based on Equation (2.2); it adapts readily to the general commutative case.

Example 2.1. In Figure 2.1(a) we see a gain graph Φ of order $n = 3$ with gains in $\mathbb{Q}^*$, the multiplicative group of rational numbers. We adopt the simplified notation ge_{ij} for an edge v_iv_j with gain $\varphi(ge_{ij}; v_i, v_j) = g$. (Then for instance $2e_{13} = 2^{-1}e_{31}$.) The balanced circles are $C_1 = \{3e_{12}, 1e_{23}, 3e_{13}\}$ and $C_2 = \{1e_{12}, 2e_{23}, 2e_{31}\}$, since their gains are $\varphi(C_1) = 3 \cdot 1 \cdot 3^{-1} = 1$ and $\varphi(C_2) = 1 \cdot 2 \cdot 2^{-1} = 1$. (The digon $\{3e_{12}, 3e_{21}\}$ has gain $3 \cdot 3 \neq 1$ so it is unbalanced.) Therefore, $\langle\Phi\rangle = (\|\Phi\|, \{C_1, C_2\})$.

The vectors that canonically represent $G(\Phi)$ in $\mathbb{R}^3$,

$$\begin{gathered}
x(h_1) = b_1 = (1,0,0), \quad x((-1)e_{22}) = -b_2 + (-1)b_2 = (0,-2,0),\\
x(1e_{12}) = (-1,1,0), \quad x(3e_{12}) = (-1,3,0), \quad x(3e_{21}) = (3,-1,0),\\
x(1e_{23}) = (0,-1,1), \quad x(2e_{23}) = (0,-1,2),\\
x(2e_{13}) = (-1,0,2), \quad x(3e_{13}) = (-1,0,3),
\end{gathered}$$

are shown in Figure 2.1(b), and the corresponding canonical hyperplane representation (see Corollary 2.2) is in Figure 2.1(c). The vectors span $\mathbb{R}^3$ since $\operatorname{rk} G(\Phi) = 3$. An incidence matrix of Φ (with edges in the same order as above) is

$$\mathbf{I}(\Phi) = \begin{bmatrix} 1 & 0 & -1 & -1 & 3 & 0 & 0 & -1 & -1 \\ 0 & -2 & 1 & 3 & -1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 2 & 2 & 3 \end{bmatrix}.$$

The gain group, $\mathfrak{G} = \mathbb{Q}^*$, embeds in several ways as a subgroup of $\mathbb{R}^*$. Each different embedding gives a different canonical bias representation of Φ and none of these are projectively equivalent (by Proposition 2.9). Still other canonical bias representations of $G(\Phi)$ may result from other gains for $\langle\Phi\rangle$, switching inequivalent to Φ, in gain group $\mathbb{R}^*$. Indeed,

(a)

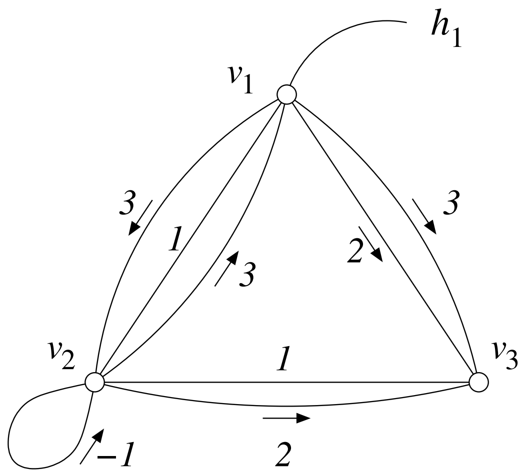

h_1
v_1
v_2
v_3
3
1
3
2
3
1
−1
2

(b)

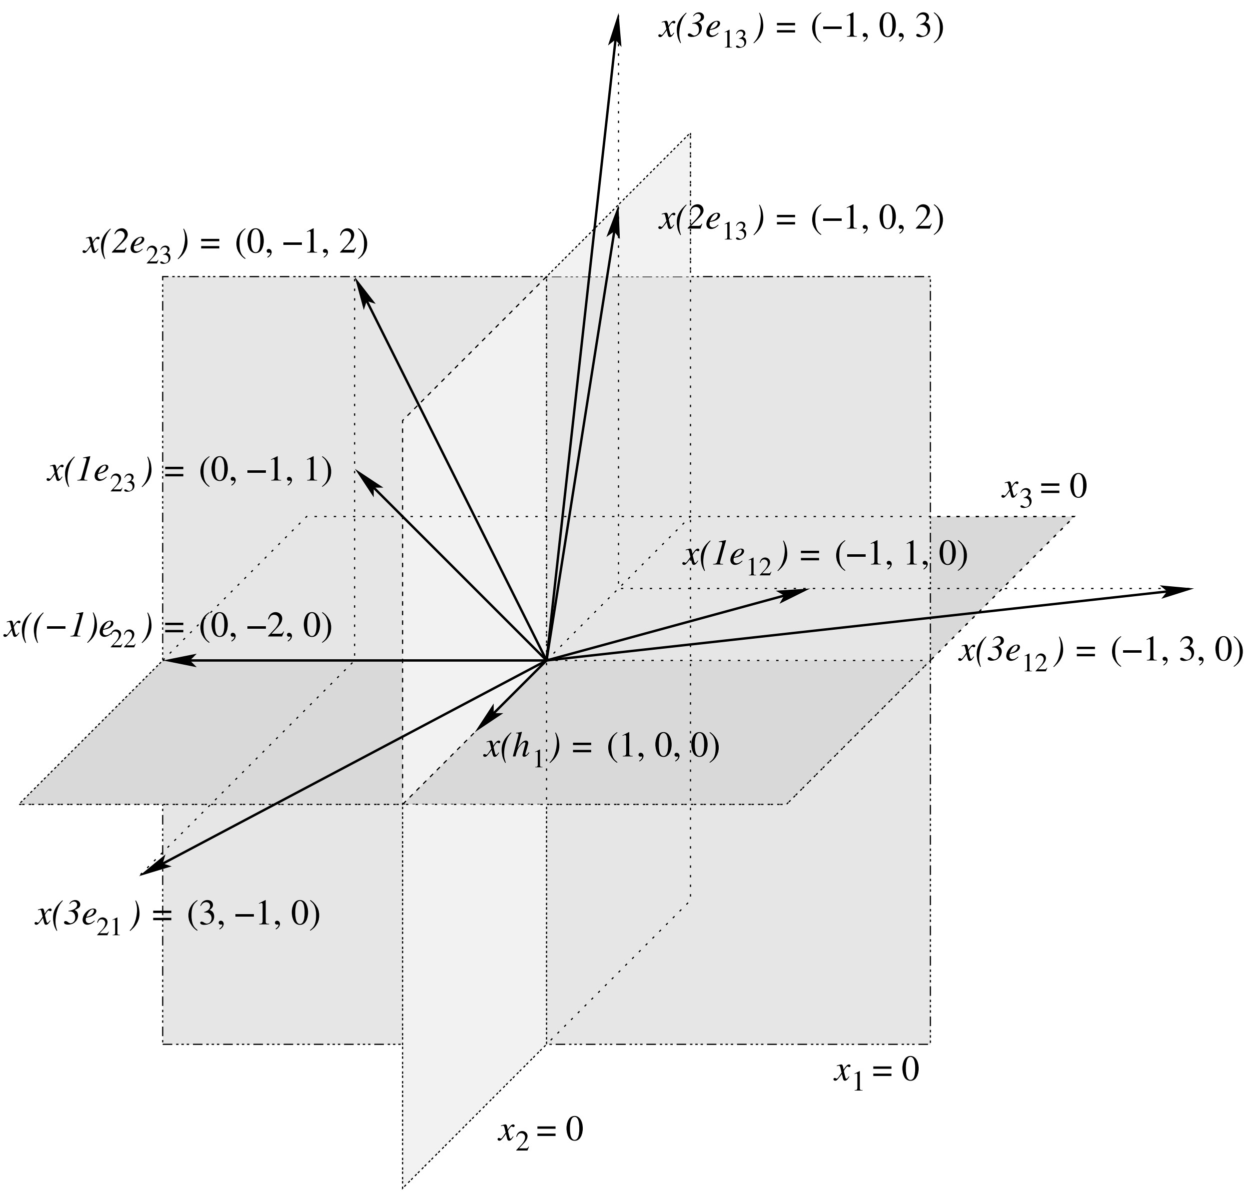

$x(3e_{13}) = (-1, 0, 3)$
$x(2e_{13}) = (-1, 0, 2)$
$x(2e_{23}) = (0, -1, 2)$
$x(1e_{23}) = (0, -1, 1)$
$x_3 = 0$
$x(1e_{12}) = (-1, 1, 0)$
$x((-1)e_{22}) = (0, -2, 0)$
$x(3e_{12}) = (-1, 3, 0)$
$x(h_1) = (1, 0, 0)$
$x(3e_{21}) = (3, -1, 0)$
$x_1 = 0$
$x_2 = 0$

(c)

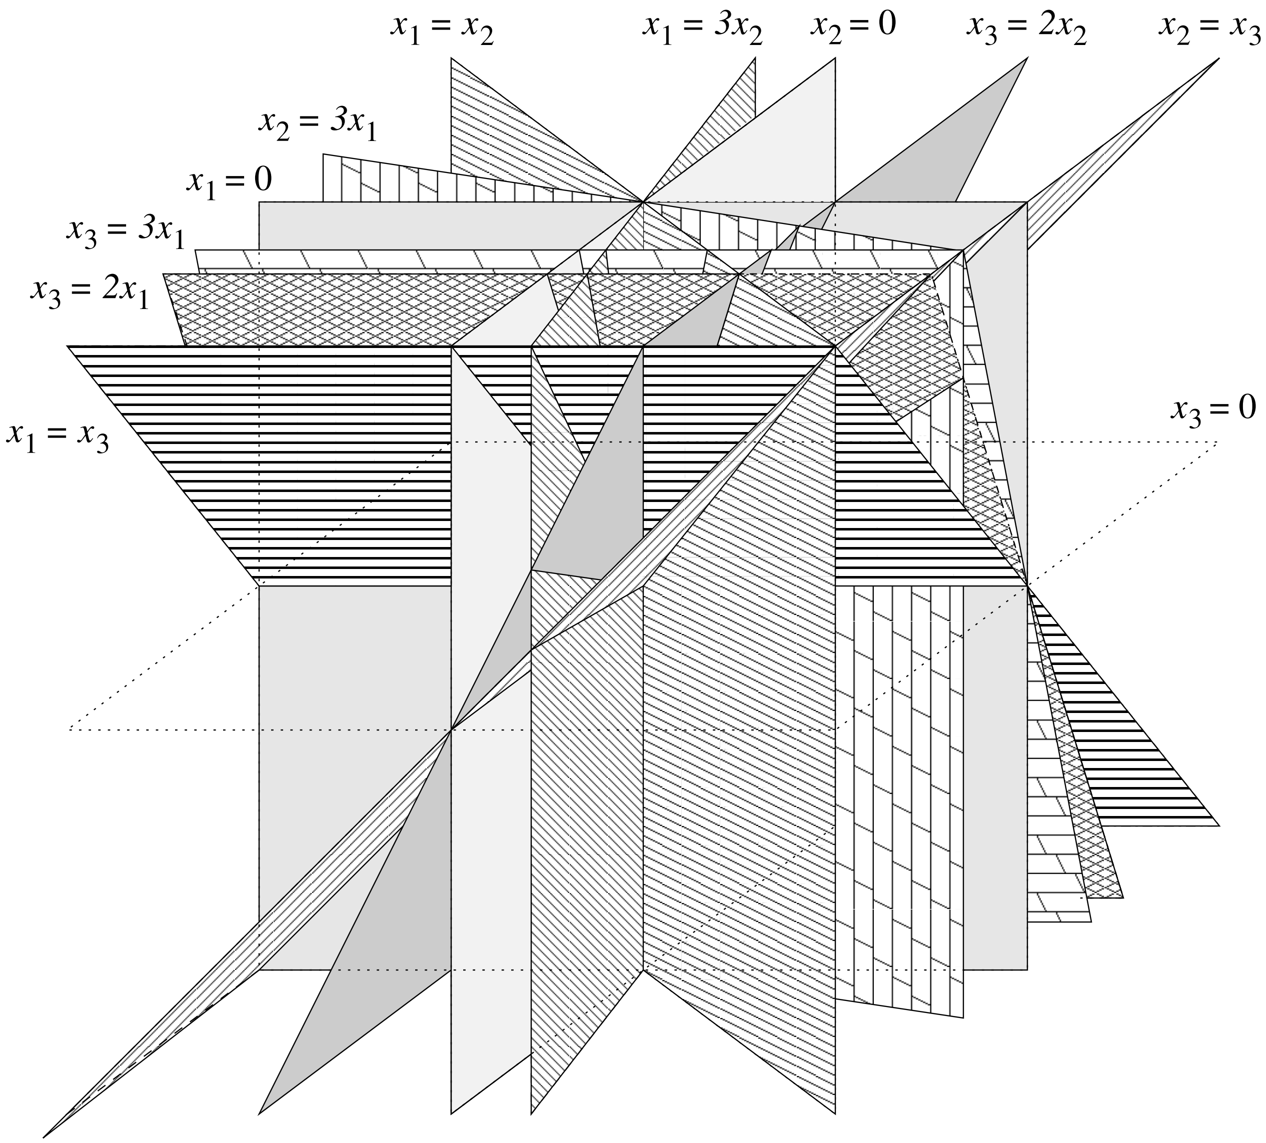

FIGURE 2.1. A $\mathbb{Q}^*$-gain graph Φ, a canonical real vector bias representation (b), and a canonical real hyperplanar representation (c). Arrows by the edges of Φ indicate the direction in which the gain is calculated; the superfluous arrows by identity-gain edges are omitted.

up to switching, gains for $\langle\Phi\rangle$ are any of these:

$$
\begin{array}{lll}
\varphi(1e_{12}; v_1v_2) = 1, & \varphi(3e_{12}; v_1, v_2) = \alpha, & \varphi(3e_{21}; v_1, v_2) = \beta, \\
\varphi(1e_{23}; v_2, v_3) = \gamma, & \varphi(2e_{23}; v_2, v_3) = 1, & \\
\varphi(2e_{13}; v_1, v_3) = 1, & \varphi(3e_{13}; v_1, v_3) = \alpha\gamma, & \\
\varphi((-1)e_{22}) = \delta, & &
\end{array}
$$

where $\alpha, \beta, \gamma, \delta \neq 1$, $\alpha \neq \beta$, $\alpha\gamma \neq 1$, $\alpha\beta \neq 1$. Even ignoring δ, which does not materially affect the representation, this leaves three nearly independent choices $\alpha, \beta, \gamma \in \mathbb{R}^*$ leading to different, projectively inequivalent canonical bias representations.

However, since every link is multiple and $G(\Phi) \neq L(\Phi)$, by Theorem 7.1 every real representation of $G(\Phi)$ is a canonical bias representation with respect to some $\mathbb{R}^*$-gains for $\langle\Phi\rangle$. We may therefore say that the inequivalent real representations of $G(\Phi)$ form a 3-parameter family. Similar remarks apply for any choice of scalar skew field F, keeping in mind that F must be large enough for α, β, γ to be choosable, that is, $\#F \geq 5$. One must also remember that parameters (α, β, γ) and $(\alpha', \beta', \gamma')$ give projectively equivalent representations if they are related by $\operatorname{Aut} F$; and only if, because of Proposition 2.9.

Example 2.2 (*Swirls*). A *swirl* is a matroid $G(2C_n, \mathcal{B})$ where $n \geq 3$ and $\mathcal{B}$ is a linear class of Hamiltonian circles. (This is a special case of Hamiltonian bias, Examples I.6.8 and III.12.2. By the definition of a linear class, no two circles in $\mathcal{B}$ can differ in only one edge.) Swirls (along with spikes, Example 4.2) are crucial examples for representability over finite fields [15, 30, 35].

A *free swirl* is a bicircular matroid $G(2C_n, \varnothing)$ with $n \geq 3$. Arbitrary swirls differ from free ones in having circuit hyperplanes corresponding to the balanced circles. From amongst the great variety of swirls [35] singles out three: the free swirls (see [30, Section 5], whose matroids M_r are free swirls), swirls with one balanced circle, and those with two complementary balanced circles. Each has its own representabilities, which we shall not go into here; we content ourselves with two observations.

First, we describe representations. The free swirl $G(2C_n, \varnothing)$, like any bicircular matroid, is representable over every sufficiently large field. Here is how: Choose n independent base points $p_1, \ldots, p_n$ (in projective space, say) and two points x_i, y_i on each line $\operatorname{span}\{p_{i-1}, p_i\}$ (with $p_0 = p_n$) so that no n of $x_1, y_1, \ldots, x_n, y_n$ are dependent; these represent the free swirl. Other swirls are obtained by specializing the x_i and y_i to form the required circuit hyperplanes while maintaining the independence of the base points and the unbalanced circles. (This is the definition of a swirl in [35].) Theorem 7.1 tells us there are no other ways to represent a swirl if $n \geq 4$. (We shall say a little more about sufficient largeness in Part VII.)

Second, the swirl $G(2C_n, \mathcal{B})$ is F-representable if and only if $(2C_n, \mathcal{B})$ has gains in F^*. Also, if $(2C_n, \mathcal{B})$ happens to have no gains in any group, then its swirl $G(2C_n, \mathcal{B})$ is not representable over any skew field. All this is also a consequence of Theorem 7.1. One example of a nongainable biased $2C_n$ with $n = 4$ is given in Example I.5.8; it yields a nonrepresentable swirl of rank 4.

There is a dual representation by hyperplanes. In F^N let x_v be the coordinate corresponding to node v. Let $h(e)$ be the hyperplane specified by $x_v = \varphi(e; v, w)x_w$ if e has endpoints v and w, by $x_v = 0$ if e is a half edge at v, and let $h(e) = F^N$ (the *'degenerate hyperplane'*) if e is a loose edge or balanced loop. Let $\mathcal{H}(\Phi)$ be the set (or multiset) $\{h(e) : e \in E\}$. We call this the *canonical linear hyperplane bias representation* of Φ.

Corollary 2.2. *The set $\mathcal{L}(\mathcal{H}(\Phi))$ of flats of $\mathcal{H}(\Phi)$, ordered by reverse inclusion, is isomorphic to* $\operatorname{Lat} G(\Phi)$ *under the mutually inverse mappings*

$$A \in \operatorname{Lat} G(\Phi) \mapsto \bigcap_{e \in A} h(e), \qquad t \in \mathcal{L} \mapsto \{e : h(e) \supseteq t\}.$$

For $S \subseteq E$, we have $\mathcal{L}(\mathcal{H})(\Phi|S)) \cong \operatorname{Lat} G(\Phi|S)$, and if $s = \bigcap_{e \in S} h(e)$, then $\mathcal{H}(\Phi/S) \cong \{h(e) \cap s : e \notin S\}$ and $\mathcal{L}(\mathcal{H}(\Phi/S)) \cong \{t \in \mathcal{L}(\mathcal{H}(\Phi)) : t \subseteq s\} \cong \operatorname{Lat} G(\Phi/S)$. $\square$

Corollary 2.3. *Let Φ be finite.*

(a) *Suppose $F = \mathbb{R}$. Then the set $x(E)$ has $\frac{1}{2}|\chi_\Phi(-1)|$ hyperplane separations. $\mathcal{H}(\Phi)$ has $(-1)^n \chi_\Phi(-1)$ regions; its face generating polynomial equals $(-1)^n w_\Phi(-x,-1)$.*

(b) *Suppose $F = \mathbb{C}$. Then the complement $\mathbb{C}^n \setminus \bigcup \mathcal{H}(\Phi)$ has Poincaré polynomial equal to $(-y)^n \chi_\Phi(-1/y)$.*

Proof. (a) By Theorem 2.1(a) and Corollary 2.2, respectively, theorems of [36] cited in Section 1, and Theorem III.5.1.

(b) By Corollary 2.2, [28], and Theorem III.5.1. $\square$

Several real and complex hyperplanar representations related to root systems and complex generalizations, with enumerative formulas, are given in Section V.8. Here we present a different kind of example (that will show up again in Menelæan and Cevian versions in Examples 2.10 and especially 2.12).

Example 2.3 (*Some real and complex hyperplanes*). Think of the canonical hyperplanar bias representations of $-K_n$ and $-K_n^\bullet$. The hyperplanes are $x_i + x_j = 0$ in the former and in the latter also $x_i = 0$. The chromatic polynomials (from [38, Equations (1.1) and (5.7)]) are

$$\begin{aligned}
\chi_{-K_n^\bullet}(\lambda) &= \sum_{i=1}^{n} S(n,i) 2^i \left(\frac{\lambda-1}{2}\right)_i, \\
\chi_{-K_n}(\lambda) &= \chi_{-K_n^\bullet}(\lambda) + n\chi_{-K_{n-1}^\bullet}(\lambda) \\
&= \sum_{i=1}^{n} \big[S(n,i) + nS(n-1,i)\big] 2^i \left(\frac{\lambda-1}{2}\right)_i,
\end{aligned}$$

where $(\lambda)_n$ is the falling factorial $\lambda(\lambda-1)\cdots(\lambda-n+1)$ and $S(n,i)$ is the Stirling number of the second kind, the number of partitions of an n-element set into i parts; so the numbers of regions of the corresponding real arrangements are

$$\begin{aligned}
r(\mathcal{H}_\mathbb{R}(-K_n^\bullet)) &= \sum_{i=1}^{d+1} S(d+1,i) 2^i i! (-1)^{d+1-i}, \\
r(\mathcal{H}_\mathbb{R}(-K_n)) &= \sum_{i=1}^{d+1} \big[S(d+1,i) + (d+1)S(d,i)\big] 2^i i! (-1)^{d+1-i}.
\end{aligned}$$

From the Poincaré polynomials given by Corollary 2.3(b) we obtain the Betti numbers of the complements

$$X = \mathbb{C}^n \setminus \bigcup \mathcal{H}_\mathbb{C}(-K_n) \quad \text{and} \quad X^\bullet = \mathbb{C}^n \setminus \bigcup \mathcal{H}_\mathbb{C}(-K_n^\bullet)$$

of the complex hyperplane arrangements: they are

$$\begin{aligned}
\beta_k(X^\bullet) &= \sum_{i=0}^{k} (-1)^i \binom{n-i}{n-k} t(n, n-i), \\
\beta_k(X) &= \sum_{i=0}^{k} (-1)^i \binom{n-i}{n-k} \left[t(n,n-i) + nt(n-1,n-i)\right],
\end{aligned}$$

where

$$t(n, n-i) = \sum_{j=0}^{i} 2^{i-j} S(n, n-j) s(n-j, n-i);$$

here $s(n,k)$ is the Stirling number of the first kind, or $(-1)^{n-k}$ times the number of permutations of $\{1, 2, \ldots, n\}$ having k cycles.

One might hope that, conversely to Theorem 2.1(a), any linear representation of the bias matroid of a biased graph Ω is a canonical bias representation of a gain graph whose biased graph is Ω. But this is not always so.

Example 2.4. Consider $G(\pm K_3)$, where $\pm K_3$ is the signed expansion of K_3; that is, it has all possible positive and negative links. By [37, Section 5] $G(\pm K_3) \cong G(K_4)$. Although $G(K_4)$ is binary, $\langle \pm K_3 \rangle$ has no canonical bias representation over $\mathbb{F}_2$. The reason is that $\langle \pm K_3 \rangle$ can only have a gain group in which some element has order 2 and there is no such element in $\mathbb{F}_2^*$. See Example 7.5 for more about this example and for representation diagrams.

There is, however, a partial converse to the theorem.

Proposition 2.4. *Let $\Omega = (\Gamma, \mathcal{B})$ be a full biased graph and F a skew field. Let f be an F-linear representation of $G(\Omega)$. Then there is a gain graph Φ such that $\langle \Phi \rangle = \Omega$ and f is a canonical bias representation of Φ.*

Proof. Suppose $f : E \to \Lambda$ is the representation of $G(\Omega)$. Since a set E_N consisting of one unbalanced edge e_v at each node v is a basis for $G(\Omega)$, $f(E_N)$ is independent. Let us fix E_N and set $\hat{v} = f(e_v)$. Then for an edge e with $\nu_\Gamma(e) = \{v, w\}$, $f(e)$ lies in span$\{\hat{v}, \hat{w}\}$ so it is a scalar multiple of $\hat{v} - \alpha\hat{w}$ for some $\alpha \neq 0$. We take $\varphi(e; v, w) = \alpha$. Now let $\Phi = (\Gamma, \varphi, F^*)$. Since f is a canonical bias representation of Φ and a representation of $G(\Omega)$, the identity mapping on E is an isomorphism of $G(\Phi)$ with $G(\Omega)$. It follows that $\langle \Phi \rangle = \Omega$. □

Proposition 2.4 generalizes the vectorial portion of the unique representation theorem for Dowling lattices [13, Theorem 9]. Dowling proved that, for $n \geq 3$ and $\#\mathfrak{G} \geq 2$, any vector representation f of $G(\mathfrak{G}K_n^\bullet)$ is a canonical bias representation. We can deduce this through the following argument. (1) By Proposition 2.4, f is a canonical bias representation of some gain graph Ψ with $\langle \Psi \rangle = \langle \mathfrak{G}K_n^\bullet \rangle$. (2) If $\langle \Psi \rangle = \langle \mathfrak{G}K_n^\bullet \rangle$, where $n \geq 3$, then Ψ is switching equivalent to $\mathfrak{G}K_n^\bullet$. (We must assume Ψ has no isolated or monovalent nodes and we must treat all unbalanced edges as if they were half edges. The gain group of Ψ may be larger than $\mathfrak{G}$ but it can be cut down after switching.) This is easily proved (or see Theorem V.2.1). (3) Since Ψ and $\mathfrak{G}K_n^\bullet$ are switching equivalent, their canonical bias representations are projectively equivalent (see the remark near Equation (2.1)). Combining (1)–(3) yields Dowling's theorem.

Define a *canonical bias representation (over F) of the biased graph* Ω to be any canonical bias representation over F of a gain graph Φ whose biased graph is Ω. (This is by definition a vector representation.) Different choices of Φ, although they have the same biased graph, need not have any other relation to each other; e.g., they may have different gain groups. By Proposition 2.4, we could equivalently define a canonical bias representation of Ω as the restriction to E of any linear representation of $G(\Omega^\bullet)$.

Thus, we define a *canonical projective* [or, *affine*] *bias representation of a biased graph* Ω to be the restriction to $E(\Omega)$ of any projective [affine] representation of $G(\Omega^\bullet)$. This definition is compatible with that of a canonical linear bias representation of Ω. To see why, suppose $G(\Omega)$

represented in a projective space $\mathbb{P}$ that has coordinates in F. $\mathbb{P}$ is therefore a quotient of an F-vector space Λ: to be precise, the points of $\mathbb{P}$ are the lines of Λ. If we pull back the point $\hat{e}$ that represents $e \in E^{\bullet}$ to a nonzero vector $f(e)$ in the line corresponding to $\hat{e}$, then f will be a canonical bias representation of some $\Phi^{\bullet}$ for which $\langle \Phi^{\bullet} \rangle = \Omega^{\bullet}$. Restricting f to E gives a linear canonical bias representation of Ω. That is, coordinatizable canonical projective representations are equivalent to linear canonical representations. If we were content to have canonical bias representations only in coordinatizable projective spaces, we could have defined a projective bias representation of Ω to be the projective image of a linear bias representation. However, that would have been unnecessarily confining. Noncoordinatizable projective representations (by Desargues' theorem, necessarily in rank ≤ 3) are in some ways just as good as coordinatizable ones—an example is in the proof of Theorem 7.1, where we can apply the modular law instead of relying on coordinates—and they give added power that becomes interesting in connection with quasigroup expansions and nondesarguesian representations in Parts V and VI.

Example 2.5 (*Balanced*). If Φ is balanced (i.e., if $\Phi = \langle \Gamma \rangle$ so for instance $\Phi = \{1\}\Gamma$), then a canonical bias representation is any representation, because $G(\Phi) = G(\Gamma)$, which has projectively unique representation.

Example 2.5A (Near-regular). The near-regular and $\sqrt[6]{1}$-matroid of maximum size of [30a], T_r, is gain graphic. Let Φ_r consist of K_{r+1} with identity gains, on node set $\{v_0, v_1, \ldots, v_r\}$, and additional edges f_{0i} with gain $\varphi(f_{0i}) = \alpha \neq 1$ for $i = 1, \ldots, r$. It follows from the description of T_r in [30a, bottom of p. 166] that $T_r = G(\Phi_r^{\bullet})$, no matter the value of α. Hence, T_r is (canonically) representable over every skew field of order at least 3. By Proposition 2.4 and Eq. (2.1) there is one representation (up to projective equivalence) for each automorphism-equivalence class of elements of $F \backslash \{0, 1\}$, and all these different representations are projectively inequivalent.

Interestingly noncanonical representations of bias matroids of biased graphs seem to be hard to find and I hardly know any. One noncanonical bias representation is the binary representation of $\pm K_3$ just mentioned in Example 2.4; but this is of the type where any two unbalanced circles have a common node so that $G(\Phi) = L(\Phi)$; thus the noncanonical bias representation is a canonical lift representation. For another example of the same kind:

Example 2.6. By [42] $G(-K_5) = R_{10}$, Bixby's regular matroid, so it has a binary representation, but this cannot be a canonical bias representation for the same reason as applies to $\pm K_3$: the only possible gains (up to switching) are all negative, and $\mathbb{Z}_2 \not\leq \mathbb{F}_2^*$. However, this is a canonical lift representation. Since $\mathbb{Z}_2 = \mathbb{F}_2^+$, $L(-K_5)$ has a canonical binary representation. A regular matroid has projectively unique representation, so the binary representation of $G(-K_5)$ can only be the canonical representation of $L(-K_5)$.

Example 2.7 (*Too thick*). A more interesting way to get noncanonical bias representations is to make the space too narrow. Suppose Ω has a k-fold link with no balanced digons. Then in $G(\Omega)$ there is a k-point line that represents only links. In $G(\Omega^{\bullet})$ this line will have $k+2$ points. $G(\Omega)$ might be representable over $\mathbb{F}_q$ where $q = k-1$ or k (if these are prime powers), but the representation is noncanonical for bias because it cannot extend to $G(\Omega^{\bullet})$. Such an example is Ω_k, a tree with one edge thickened to k parallel edges and with no balanced digons. Then $G(\Omega_k) = U_{2,k} \oplus U_{n-2,n-2}$, the direct sum of a line and a free matroid, which is obviously $\mathbb{F}_q$-representable for any $q \geq k-1$.

This kind of noncanonicality is still rather superficial, in that the representation becomes canonical upon extending the coordinate field.

Problem 2.5 (*Fundamental representation questions*). (a) Characterize the biased graphs for which the conclusion of Proposition 2.4 does not hold: that is, they have a noncanonical bias representation (not counting canonical lift representations when $G(\Omega) = L(\Omega)$); most especially, one that is not binary. (b) What kinds of noncanonical representation can exist in these cases?

2.2. **Right canonical representations.**

The previous work is valid if F acts on Λ on the left. If the action is on the right (and F is noncommutative), certain modifications are necessary.

Define $x'_\Phi(e) = \hat{v}\varphi(e; v, w) - \hat{w}$. We call a mapping $f : E \to \Lambda$ a *right canonical bias representation of* Φ if there exists a mapping $N \to \Lambda$ of N onto an independent set under which each $f(e)$ is a non-zero multiple of $x'_\Phi(e)$. (The previous representations would in this context be called *left canonical.* When F is commutative, left and right canonical representations are not the same.)

Theorem 2.6. *Let $\Phi = (\Gamma, \varphi, \mathfrak{G})$ be a gain graph and F a skew field containing $\mathfrak{G}$ as a multiplicative subgroup. Then the right-linear dependence matroid of the vectors $x'_\Phi(e)$ is naturally isomorphic to the bias matroid $G(\Phi)$.*

First Proof. Consider a walk $W = (v_0, e_1, v_1, e_2, \ldots, e_l, v_l)$ with gain $\varphi(W) = \varphi(e_1)\varphi(e_2)\cdots\varphi(e_l)$. It is clear that

$$\sum_{i=1}^{l} x'(e_i)\alpha_i = \hat{v}_0\varphi(W) - \hat{v}_l$$

if we take $\alpha_i = \varphi(e_{i+1})\cdots\varphi(e_l)$. Thus a proof like that of Theorem 2.1(a) goes through. □

Second Proof. The left and right canonical representations differ when F is a field. The true relationship is rather complex. First, note that $-x'_\Phi(e)\varphi(e)^{-1} = -\hat{v} + \hat{w}\varphi(e; v, w)^{-1}$. This suggests an equivalence between left and right representations through $\mathfrak{G}^{\mathrm{op}}$ and F^{op}, the opposite group and skew field with multiplications $\circ$ defined by $g \circ h = h \cdot g$, where the dot is multiplication in $\mathfrak{G}$ or F. The right F-vector space Λ has a corresponding left F^{op}-vector space Λ^{op}. Let $\Phi^{\mathrm{op}} = (\|\Phi\|, \varphi^{\mathrm{op}}, \mathfrak{G}^{\mathrm{op}})$ where $\varphi^{\mathrm{op}}(e) = \varphi(e)^{-1} \in \mathfrak{G}^{\mathrm{op}}$. Then $\varphi^{\mathrm{op}}(e_1 e_2 \cdots e_l) = \varphi(e_1 e_2 \cdots e_l)^{-1}$, so Φ^{op} has the same balance as Φ, whence $G(\Phi^{\mathrm{op}}) = G(\Phi)$. Now,

$$\begin{aligned}[-x'_\Phi(e)\varphi(e)^{-1}]^{\mathrm{op}} &= -\hat{v} + \varphi(e; v, w)^{-1}\hat{w} \\ &= -\hat{v} + \varphi^{\mathrm{op}}(e; v, w)\hat{w} = x_{\Phi^{\mathrm{op}}}(e).\end{aligned}$$

Thus the natural reversal mapping $\Lambda \to \Lambda^{\mathrm{op}}$ maps a scalar multiple of $x'_\Phi(e)$ to $x_{\Phi^{\mathrm{op}}}(e)$. This establishes that x'_Φ and $x_{\Phi^{\mathrm{op}}}$ represent the same matroid; by Theorem 2.1(a) it is $G(\Phi^{\mathrm{op}})$, which equals $G(\Phi)$. □

Right canonical representations and the opposite gain graph give us extra flexibility in defining hyperplane representations. Regarding F^N as a right F-vector space, let $h'(e)$ be the hyperplane specified by $x_w = x_v\varphi(e; v, w)$. I usually prefer h' to h as a hyperplane representation of $G(\Phi)$,[1] but it needs justification. Let $\mathcal{H}'(\Phi) = \{h'(e) : e \in E\}$. The analog of Corollary 2.2 is

Corollary 2.7. *Corollary 2.2 remains valid if $\mathcal{H}$ and h are replaced by $\mathcal{H}'$ and h'.*

Proof. First proof: h' is dual to x'. Second proof: apply Φ^{op} as in the second proof of Theorem 2.6. □

2.3. **Switching and projective equivalence.**

At Equation (2.1) we noted that the canonical bias representations of members of a switching class are projectively equivalent. The converse is not true for two reasons. Suppose Φ is an F^*-gain graph. First of all, we need to allow both switching and field automorphisms: for a switching function η and an $\alpha \in \operatorname{Aut} F$, and writing $\approx$ for projective equivalence, we have $x_{\Phi^{\eta\alpha}} \approx x_\Phi$. Secondly, sometimes Φ' and Φ on the same underlying graph can have $x_{\Phi'} \approx x_\Phi$ (so $G(\Phi') = G(\Phi)$, whence $\langle\Phi'\rangle = \langle\Phi\rangle$) but $\Phi' \neq \Phi^{\eta\alpha}$; this is certainly possible for a contrabalanced circle or theta graph. Those may be essentially the only such examples.

Conjecture 2.8. Suppose Φ and Φ' are unbalanced F^*-gain graphs of finite order with $\|\Phi'\| = \|\Phi\|$ and with $G(\Phi)$ connected. If $G(\Phi)$ has a $U_{2,4}$ minor, and if $x_{\Phi'} \approx x_\Phi$, then Φ' is obtained from Φ by switching and a field automorphism. If $G(\Phi)$ has no $U_{2,4}$ minor and $x_{\Phi'} \approx x_\Phi$, then Φ' need not be so obtained.

I propose the sufficiency of a $U_{2,4}$ minor because a quadruple of points on a projective line coordinatized by F is projective to another such quadruple if and only if they have the same cross ratio, and the matroid of a quadruple is $U_{2,4}$. (The cross ratio in a skew field is a conjugacy class; see [3, p. 72, Proposition 1].) This leads me to think that it is sufficient for the conjecture if every element of $G(\Phi)$ belongs to a $U_{2,4}$ minor. By a result of Bixby [4], if $G(\Phi)$ is connected and has a $U_{2,4}$ minor, every element belongs to a $U_{2,4}$ minor; therefore we may replace the hypothesis of multielement $U_{2,4}$'s by the assumption that $G(\Phi)$ is connected (this is characterized in Theorem II.2.8) and nonbinary (characterized in [40, Theorem 3]).

2.4. **Abstract gains and nonunique representation.**

We should emphasize that in Theorem 2.1 $\mathfrak{G}$ is a specific subgroup of F^* leading to a specific canonical bias representation. If $\mathfrak{G}$ is known only as an abstract group, it may embed as a subgroup of F^* in several ways that are not equivalent under automorphisms of F.[2] Since the question of projectively unique representability of a matroid is an important one, we need a theorem to tell us how that uniqueness is affected by existence of multiple embeddings of the gain group.

To that end we define a *canonical bias representation* of a gain graph Φ with abstract gain group $\mathfrak{G}$ to be a canonical bias representation of Φ over F, as in Theorem 2.1, obtained by choosing any embedding $\varepsilon : \mathfrak{G} \hookrightarrow F^*$.

Proposition 2.9. *Suppose Φ is a loopless gain graph whose gain group $\mathfrak{G}$ is generated by $\varphi(E_*)$. The canonical bias representations of Φ induced by different embeddings $\varepsilon_1, \varepsilon_2 : \mathfrak{G} \hookrightarrow$*

[1] The choice between h and h' in work with hyperplane representations is so arbitrary that one should just choose one and then forget the special notation.

[2] For instance, $\mathbb{Z} \hookrightarrow \mathbb{Q}^*$ in infinitely many ways, but $\mathbb{Q}$ has no nontrivial automorphisms. A different kind of example is where F has automorphisms but $\mathfrak{G}$ has more; thus when $\gamma \geq 5$, $\mathbb{Z}_\gamma \hookrightarrow \mathbb{C}^*$ in inequivalent ways. One more example: if $k > 2$ and $p \geq 5$, then $\mathbb{Z}_k \hookrightarrow \mathbb{F}_{p^d}^*$ both by $1 \mapsto 2$ and by $-1 \mapsto 2$ but 2 and 2^{-1} are not automorphic.

F^* *are projectively equivalent if and only if* ε_1 *and* ε_2 *are equivalent under an automorphism of* F.

Proof. Let f_i be a canonical bias representation in F^N induced by ε_i: that is, with respect to some basis $\hat{N}_i$, $f_i(e)$ is a scalar multiple of $x_i(e) = -\hat{v}_i + \varepsilon_i(\varphi(e))\hat{w}_i$. By linear transformation of F^N we may assume $\hat{v}_1 = \hat{v}_2$ for all $v \in N$. By scaling we may assume $f_i = x_i$. Consequently,

$$x_2(e) = x_1(e)^\alpha = -\hat{v} + \varepsilon_1(\varphi(e))^\alpha \hat{w}$$

for some automorphism α of F, from which it follows that $\varepsilon_2(\varphi(e)) = \varepsilon_1(\varphi(e))^\alpha$ (since $v \neq w$). We conclude that ε_2 is the composition of ε_1 and α. □

In light of this result, unique representation has multiple levels of interpretation. One interpretation takes account of the gain-group embedding. We say that Φ has *projectively unique bias representation up to gain-group embedding* if every bias representation is projectively equivalent to a canonical bias representation with respect to some embedding of $\mathfrak{G}$ in F^*.

2.5. **Menelaus.**

A third pair of dual representations of the bias matroid generalizes the theorems of Menelaus and Ceva[3] to arbitrarily high dimensions and complicated configurations. Let Φ be a gain graph with gain group contained in F^*, where F is a field. (We may as well take the gain group to be F^* itself.) Let h_∞ be the ideal hyperplane in $\mathbb{P} = \mathbb{P}^{N-1}$, so that $\mathbb{P} \setminus h_\infty$ is the affine space $\mathbb{A} = \mathbb{A}^{N-1}$. Choose an affine basis $N_\mathbb{P} = \{v_\mathbb{P} : v \in N\}$ (which is also a projective basis), define $l_{vw} = \operatorname{span}\{v_\mathbb{P}, w_\mathbb{P}\}$, and let $W_\mathbb{P} = \{w_\mathbb{P} : w \in W\}$ for $W \subseteq N$. We represent $e \in E$ by a 'Menelæan' point and (Section 2.6) a 'Cevian' hyperplane. (To keep things simple we assume Φ has no loose edges or balanced loops; they would not correspond to any points.)

If e is a half edge or unbalanced loop at v, take $p(e) = p_\Phi(e) = v_\mathbb{P}$. If e is a link whose endpoints are v and w, let

$$p(e) = p_\Phi(e) = \begin{cases} \frac{1}{1-\alpha}(v_\mathbb{P} - \alpha w_\mathbb{P}) & \text{if } \alpha = \varphi(e; v, w) \neq 1, \\ h_\infty \wedge l_{vw} & \text{if } \alpha = 1. \end{cases} \tag{2.3}$$

Let $\mathcal{M}(\Phi) = \{p(e) : e \in E\}$. This is the *projective Menelæan representation* of $G(\Phi)$. If we want an affine representation we can have no identity-gain links; then $\mathcal{M}(\Phi) \subseteq \mathbb{A}$ is the *affine Menelæan representation.*

Let us play with the first line of (2.3) using arithmetic with infinity, where $1/\infty = 0$, etc. If we substitute $\alpha = 0$, then $p(e) = v_\mathbb{P}$. With $\alpha = \infty$, $p(e) = w_\mathbb{P}$. If $\alpha = 1$, $p(e)$ is infinite, which we interpret as the point $h_\infty \wedge l_{vw}$. Thus, if we define

$$\varphi(h_v; v, w) = 0 \quad \text{and} \quad \varphi(h_v; w, v) = \infty$$

for an unbalanced edge h_v at v and for $w \neq v$, we can dispense with the second line of (2.3) and the special rule for unbalanced edges, and still $p(e)$ will always be the correct Menelæan point.

Theorem 2.10 (Generalized Theorem of Menelaus)**.** *The set of flats spanned by* $\mathcal{M}(\Phi)$ *is isomorphic to* $\operatorname{Lat} G(\Phi)$ *under the natural isomorphism induced by* $e \mapsto p(e)$. *Moreover,*

[3] For which see, for example, [1, 11]; the latter was known in the eleventh century in Moorish Spain [21]

if $S \subseteq E$ and $W \subseteq N$, the flat generated by the points $p(e)$ for $e \in S$ contains $\operatorname{span}(W_{\mathbb{P}})$ *if and only if $W \subseteq N_0(S)$.*

Proof. Let $F^{1+N} = \{(x_0, x) : x_0 \in F,\ x \in F^N\}$ and re-embed F^N in F^{1+N} as the subspace $x_0 = \sum_v x_v$. (This is valid for x with finite support, which is enough for us because both $G(\Phi)$ and the projective dependence matroid of $\mathcal{M}(\Phi)$ are finitary.) We project the embedded F^N to $\mathbb{P}$ by taking homogeneous coordinates and to $\mathbb{A}$ by taking h_∞ to be where $x_0 = 0$. Let $b = (1, 0, 0, \ldots, 0) \in F^{1+N}$ and let $\bar{v}$ be the unit basis vector of F^{1+N} in the v-direction for $v \in N$. Let $\hat{v} = b + \bar{v}$ for $v \in N$. Then $\hat{v}$ projects to $v_{\mathbb{P}} \in \mathbb{P}$. Also, $\hat{N}$ is a basis for F^N in its special embedding. Now, let $x(e)$ be the standard representation of e in F^N as embedded. It projects to $p(e)$. Since the projective relations of the projected vectors are the same as the linear relations of the vectors, we have the first part of the theorem.

To prove the second part we should augment Φ to a full gain graph $\Phi^\bullet$. Let s be the flat generated by $\{p(e) : e \in S\}$. We may as well take S to be as large as possible. Since s is a flat of $\mathcal{M}(\Phi^\bullet)$ as well, let $S^\bullet = \{e \in E^\bullet : p(e) \in s\}$. Now, $p(w) \in s$ if and only if an unbalanced edge at w lies in $S^\bullet$; and that occurs if and only if $w \in N_0(S^\bullet) = N_0(S)$. □

Theorem 2.10 is an extension to more varied point configurations of the multidimensional Menelaus theorem found in [7, 10, 23]. The latter result concerns only the case in which $\|\Phi\|$ is a circle. (But this case is fundamental.)

In order to make our Menelaus theorem fully geometrical we have to know how to construct Φ from a configuration of points. Suppose we have an affine basis B and a set P of points lying on the projective lines determined by pairs of basis elements. We define Φ to have node set $N = B$ and edge set $E = P$. A point $p \in P \cap B$ becomes a half edge. Any other point $p \in l_{vw} = \operatorname{span}\{v, w\}$, where $v, w \in B$, becomes a link with endpoints v and w and with gain determined as follows. If $p \in h_\infty$, $\varphi(e) = 1$. Otherwise, if we express $p = (1-\lambda)v + \lambda w$, then $\varphi(e; v, w) = 1/(1-\lambda^{-1})$. If we express $v = p + \alpha(w - v)$ and $w = p + \beta(w - v)$, then $\beta = \alpha + 1 = 1 - \lambda$ and $\varphi(e; v, w) = \alpha/\beta$. If there is a metric in the affine space, then $\varphi(e; v, w) = d(p, v)/d(p, w)$ where we orient l_{vw} and d represents the signed distance measured along the oriented line. Now, knowing Φ, we can apply Theorem 2.10 to deduce the matroid of the configuration. (We continue this line of study in Part VI, where we make it coordinate-free, or synthetic.)

If e and f are links between v and w (and if F is a field), the cross ratio of their Menelæan points with respect to $v_{\mathbb{P}}$ and $w_{\mathbb{P}}$ has the simple expression

$$(v_{\mathbb{P}}, w_{\mathbb{P}};\ p(e), p(f)) = \varphi(e; v, w)/\varphi(f; v, w). \tag{2.4}$$

Thus $p(e)$ and $p(f)$ are harmonic conjugates with respect to $v_{\mathbb{P}}$ and $w_{\mathbb{P}}$ precisely when the edge gains sum to zero. We can interpret this fact as a harmonic relationship between Φ and $-\Phi = (\Gamma, -\varphi)$: for each link e, $p_\Phi(e)$ and $p_{-\Phi}(e)$ are harmonically conjugate with respect to the appropriate points of $N_{\mathbb{P}}$.

Example 2.8. Take a 3-simplex with vertices $v_1, \ldots, v_4$ in $\mathbb{A}^3(\mathbb{R})$. Choose some points: p_{12}, one-third of the way from v_1 to v_2; p_{13}, one-third of the way from v_1 to v_3; and p_{14}, three-fourths of the way from v_1 to v_4. The points p_{1j} determine a plane h that intersects the other lines l_{ij} at points p_{ij}. Let us determine these points. Their Menelæan gain graph Φ is balanced because they are coplanar. The gains $\varphi(e_{12}; v_1, v_2) = -\frac{1}{2}$, $\varphi(e_{13}; v_1, v_3) = -\frac{1}{2}$, $\varphi(e_{14}; v_1, v_4) = -3$, are calculated from the geometry; the others shown in Figure 2.2 are obtained by balance of Φ. Reversing the geometrical calculation we find that p_{23} is the ideal

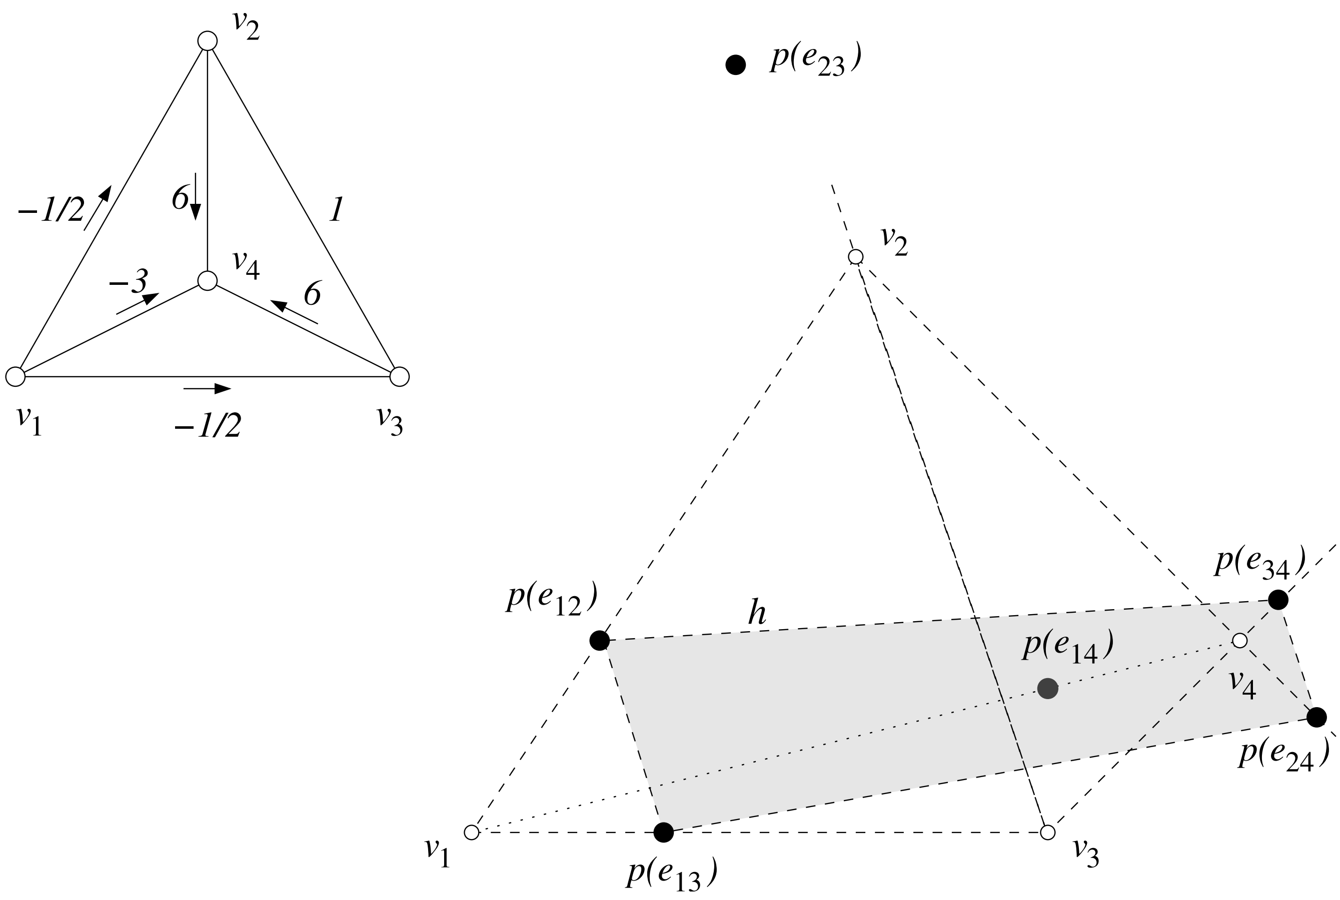

FIGURE 2.2. A real multiplicative gain graph Φ and the corresponding Menelæan points (which are coplanar, contained in h) of a tetrahedron in $\mathbb{A}^3(\mathbb{R})$, as in Example 2.8.

point on l_{23}, which means that $h \parallel l_{23}$, that p_{24} is the point on the extension of edge v_2v_4 through v_4 at distance $\frac{1}{5}d(v_2, v_4)$ past v_4, and that p_{34} is similar with respect to v_3 and v_4.

Example 2.9 (*Midpoints*). Take the midpoints of some edges in a d-dimensional simplex s^d in $\mathbb{A}^d(\mathbb{R})$. (d is $n-1$.) When are they contained in a hyperplane? For that matter, what is the dimension of the subspace they span?

Corollary 2.11 (Midpoints). *The matroid of all midpoints of a simplex s^d in real affine d-space is naturally isomorphic to $G(-K_n)$ and the matroid of midpoints and vertices of s^d is isomorphic to $G(-K_n^\bullet)$.*

Choose the midpoints of an arbitrary set S of edges. The midpoints span a flat of dimension $d-b(S)$ where $b(S)$ is the number of bipartite components of S treated as a spanning subgraph of the 1-skeleton K_{d+1} of s^d. Furthermore, the midpoints span a flat that contains no vertex if and only if S is bipartite. In particular, a maximal set of midpoints that is contained in a hyperplane that does not contain a vertex corresponds to an edge cutset of K_{d+1}.

Proof. The matroid of all midpoints and vertices is $G(-K_n^\bullet)$ by Theorem 2.10, since $-K_n$ is only a notational variant of $\{-1\}K_n$. The flat spanned by the midpoints from S has

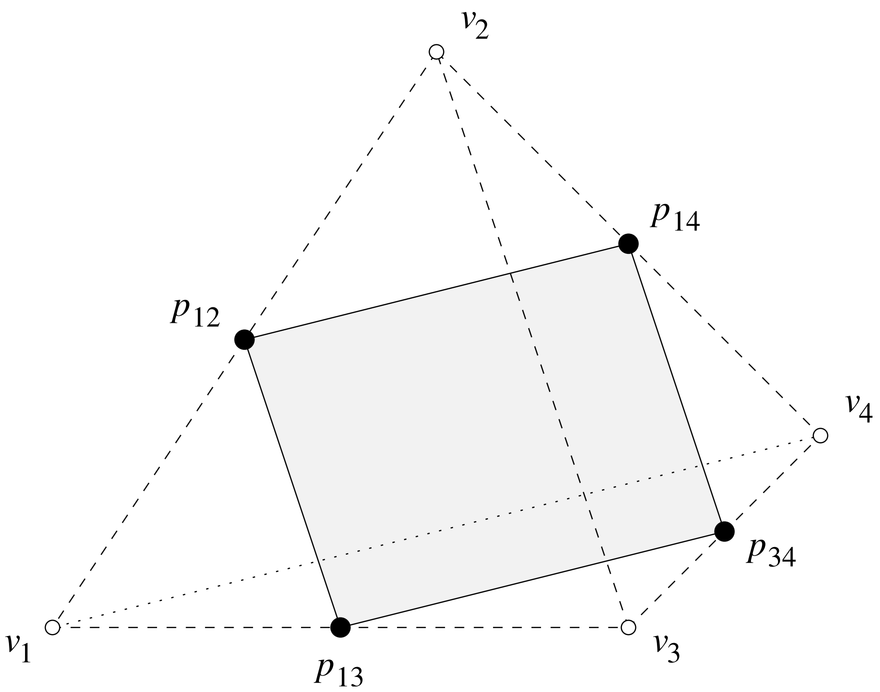

FIGURE 2.3. The midpoints of a cutset of a tetrahedron. The cutset separates v_1, v_3 from v_2, v_4.

dimension $d - b(S)$ where $b(S)$ is the number of balanced components of S in $-K_{d+1}$. A component is balanced in $-K_n$ if and only if it is bipartite. The remainder of the proof is straightforward. $\square$

The last part of the corollary generalizes the well known fact that the midpoints of the edges extending from one vertex are cohyperplanar. Figure 2.3 illustrates this. Moreover, the hyperplane is parallel to the faces of the tetrahedron that it does not intersect; this is a consequence of the more sophisticated Corollary 2.16 (as well as being clear from elementary geometry).

Regard $N_{\mathbb{P}}$ as sitting in affine space. Its convex hull is a simplex s^{N-1} of which the 1-skeleton is K_N. An arbitrary generic hyperplane h (that is, one disjoint from $N_{\mathbb{P}}$) intersects each edge line l_{vw} at a point that is either inside or outside the edge $v_{\mathbb{P}}w_{\mathbb{P}}$ of the simplex. We can easily decide which it is by examining the gain graph corresponding to the points $h \wedge l_{vw}$ for all $v \neq w \in N$. Because h is generic, h meets every edge line in a point other than a vertex of the simplex, so this gain graph is complete, having the form (K_N, φ), and it is balanced by Theorem 2.10.

Corollary 2.12. *Take a generic hyperplane h, corresponding to the gain graph $\Phi = (K_N, \varphi)$. Then h intersects l_{vw} inside the simplicial edge $v_{\mathbb{P}}w_{\mathbb{P}}$ if and only if $\varphi(e_{vw}) < 0$. The set of edges of s^{N-1} that are cut by h is either void or a cutset of K_N.*

Proof. By formula (2.3), $p(e)$ lies in an edge of the simplex if and only if $\varphi(e) < 0$. $\square$

What we have done is to create a system of coordinates in the union of all edge lines, $L = \bigcup_{v \neq w} l_{vw}$. We call these *Menelæan coordinates*: that is, the Menelæan coordinate of a point $p(e) \in L$ is $\varphi(e)$, with due regard for the special values 0 and ∞ as explained just above.[4] The gain graph associated with the coordinate system on all of L is $\Phi = F^* K_N^\bullet$. We can think of the coordinates as the composition of the mapping $p^{-1} : L \to E$ with the gain function φ of $F^* K_N^\bullet$.

There is one difficulty: p is not intrinsically geometric but is determined by the choice of homogeneous coordinate vectors $v_\mathbb{P}$ for $v \in N$. We have implicitly assumed that the $v_\mathbb{P}$ were chosen so that each one has coordinate $x_0 = 1$, as is the standard manner of embedding $\mathbb{A}^{N-1}$ in F^N. Associated with this choice is that the ideal points in L, identified by having $x_0 = 0$, are differences of base points.[5] How, then, do Menelæan coordinates change if we assign different homogeneous coordinates $v'_\mathbb{P}$ to the base points, or if we choose a different ideal hyperplane (we call this *reaffinizing* $\mathbb{P}$)? To answer these questions we need a thorough analysis of the process of setting up coordinates.

First of all, we write $\dot{v}_\mathbb{P}$ for a particular choice of homogeneous coordinate vector for $v_\mathbb{P}$, perhaps the standard one with $x_0 = 1$ or perhaps not. Since we want a coordinate system in L, which is fixed, the set $\{v_\mathbb{P}\}_{v \in N}$ itself is determined, at least unless $\#N \leq 2$, which is trivial. The only choice we can make is that of the homogeneous coordinate vectors $\dot{v}_\mathbb{P}$ for $v \in N$; that is, we can *scale* $\dot{v}_\mathbb{P}$ to $\dot{v}'_\mathbb{P} =$ a nonzero scalar multiple of $\dot{v}_\mathbb{P}$. Any such choice implies an ideal hyperplane h_∞, namely that spanned by all the vectors $\dot{w}_\mathbb{P} - \dot{v}_\mathbb{P}$. (It is easy to see that h_∞ exists and is unique.) A *Menelæan coordinate system* on L is the system of coordinates $m_{vw}(q)$ on L given by (2.3) with $m_{vw}(q) = \varphi(p^{-1}(q); v, w)$; that is, we treat q as $p(e)$, and the coordinate of q is $\varphi(e)$. (The coordinate depends on orienting l_{vw} just as the gain depends on orienting e.) Intrinsic to a Menelæan coordinate system is the bijection $p : E \to L$, which depends on the homogeneous coordinates of the base points.

Scaling the homogeneous base coordinates to $\{\dot{v}'_\mathbb{P}\}_{v \in N}$ implies a change of p to p' and hence of m to m', new Menelæan coordinates on L. We want to express m' as a change of gains in Φ. That is, we have $\varphi = p \circ m = p' \circ m'$; but if we fix the correspondence $p : E \to L$, how is $\varphi'(e) = m'(p(e))$ related to $\varphi(e) = m(p(e))$? Having fixed p, we can consider the gains φ and φ' to be the old and new Menelæan coordinates, whence the statement of our lemma:

Lemma 2.13 (Change of Menelæan coordinates). *Scaling the basis coordinate vectors $\dot{v}_\mathbb{P}$ by $\gamma : N \to F^*$, from $\dot{v}_\mathbb{P}$ to $\dot{v}'_\mathbb{P} = \gamma_v v_\mathbb{P}$, gives new Menelæan coordinates $\varphi' = \varphi^{\gamma^{-1}}$, where γ^{-1} is defined by $\gamma^{-1}(v) = \gamma(v)^{-1}$.*

Proof. A Menelæan point $q = p(e)$ given by

$$q = (1 - \lambda)\dot{v}_\mathbb{P} + \lambda \dot{w}_\mathbb{P},$$

[4] The Menelæan coordinate on a line is known in projective geometry, although not given much attention. It is the inhomogeneous parametric coordinate of [34, Section 65, p. 182]. Our definition by a distance formula, implying that $\varphi(e) = 1$ corresponds to infinity, is that of the parameter t in [24, p. 717]. (These citations are not supposed to be original sources.) What seems to be new is the use of the coordinates on a system of coordinate lines (that is, L), the systematic geometric interpretation, and the connection with gain graphs.

[5] This observation is based on the treatment of parametric coordinates of a projective line in Hughes and Piper [22, p. 33]. If we take their $e_1 = w - v$ and $e_2 = v$, then their parametric coordinates are our Menelæan coordinates; however, the geometric meaning of the coordinates is very different.

so $\varphi(e) = (1-\lambda^{-1})^{-1}$, is expressed in terms of $\dot{v}'_{\mathbb{P}}$ and $\dot{w}'_{\mathbb{P}}$ as

$$q = (1-\lambda)\gamma_v^{-1}\dot{v}'_{\mathbb{P}} + \lambda\gamma_w^{-1}\dot{w}'_{\mathbb{P}}.$$

The revised Menelæan coordinate is $\varphi'(e) = (1-\lambda'^{-1})^{-1}$ where $(1-\lambda')\dot{v}'_{\mathbb{P}} + \lambda'\dot{w}'_{\mathbb{P}}$ is a scaling of q. A short calculation shows that

$$\lambda' = [(1-\lambda)\gamma_v^{-1} + \lambda\gamma_w^{-1}]^{-1}\lambda\gamma_w^{-1} = \gamma_v[\lambda^{-1} - 1 + \gamma_w^{-1}\gamma_v]^{-1}\gamma_w^{-1},$$

so that

$$\varphi'(e) = (1-\lambda'^{-1})^{-1} = \gamma_v(1-\lambda^{-1})^{-1}\gamma_w^{-1} = \gamma_v\varphi(e)\gamma_w^{-1}.$$

Thus, $\varphi' = \varphi^{\gamma^{-1}}$. □

If γ is constant, then although φ' may not be φ, the ideal hyperplane is unchanged. This is clear from the projective geometry, but it also reflects the fact that the ideal hyperplane corresponds to the subgraph $\{1\}K_N$, whose gains are invariant. Conversely, if γ is nonconstant the gains on $\{1\}K_N$ change so the ideal hyperplane moves.

Corollary 2.14. *Assume $\#N \geq 3$ and a fixed embedding $N \to \mathbb{P}$. Any two Menelæan coordinate systems for $\bigcup_{v\neq w} l_{vw}$ are related by switching $\Phi = F^*K_N^\bullet$, and any switching of Φ gives a Menelæan coordinate system. Switching retains the ideal hyperplane in place if and only if the switching function is constant.* □

Now that we know the effect of basis scaling we can analyze reaffinization. Remember that $h_\infty \cap L$ is identified by having gains equal to 1.

Theorem 2.15 (Reaffinization). *Given: the gain graph $\Phi = F^*K_N^\bullet$, a bijection of N to an affine basis $N_{\mathbb{P}} \subseteq \mathbb{A}^{N-1}(F)$ with associated ideal hyperplane h_∞, and the corresponding Menelæan representation $p : E \to L = \bigcup_{v\neq w} l_{vw}$. Let h' be a projective hyperplane disjoint from $N_{\mathbb{P}}$ and let $H' = \{e : p(e) \in h'\}$. To reaffinize $\mathbb{P}^{N-1}(F)$ so h' becomes the ideal hyperplane, one must switch Φ to $\Phi' = \Phi^{\pi^{-1}}$ where $\pi : N \to F^*$ is any potential for $\Phi|H'$. No other switching makes h' into the ideal hyperplane.*

Proof. Because h' is a hyperplane and contains no point of $N_{\mathbb{P}}$, $\Phi|H'$ is balanced by Theorem 2.10. If π is a potential for $\Phi|H'$, then $\Phi^{\pi^{-1}}|H' = \{1\}K_N$, which is the edge set corresponding to the new ideal hyperplane h'_∞. Hence $L \cap h'$ is contained in h'_∞, and since it spans h', $h' = h'_\infty$.

It is clear that only switching by π^{-1} for a potential π can give H' all identity gains. □

The value of the theorem is that it lets us treat any hyperplane in general position with respect to $N_{\mathbb{P}}$ as the ideal hyperplane. In other words, h_∞ is a typical hyperplane.

Example 2.10 (*Midpoints and farpoints*). We shall call *farpoints* the ideal points of the edge lines of a simplex. We want a generalization of Corollary 2.11 to include the farpoints. The farpoints themselves, obviously, all belong to a hyperplane that contains no vertex; the most interesting question concerns mixtures of midpoints and farpoints.

Corollary 2.16 (Midpoints and farpoints). *The matroid of all vertices, midpoints, and farpoints of an affine simplex s^d in $\mathbb{A}^d(\mathbb{R})$ is naturally isomorphic to $G(\pm K_{d+1}^\bullet)$, with vertices corresponding to half edges, midpoints to negative edges, and farpoints to positive edges.*

A maximal set of cohyperplanar midpoints and farpoints is one of four types:

(a) *all farpoints of s^d;*

(b) *the midpoints of edges in a cutset of K_{d+1}, the 1-skeleton of s^d, and the farpoints of the edges not in the cutset;*

(c) *all midpoints and farpoints in a k-face of s^d with $0 < k < d$ and all farpoints (if any) of the opposite face;*

(d) *all midpoints and farpoints in a k-face of s^d with $0 < k \leq d-2$, all midpoints of a cutset in the 1-skeleton K_{d-k} of the opposite face, and all farpoints of edges in the opposite face but not in the cutset.*

In cases (a) and (b) the hyperplane contains no vertices. In case (b) the hyperplane is generated by its midpoints and is parallel to and equidistant from the faces of s^d that it separates.

Proof. The isomorphism is due to Theorem 2.10. The description of hyperplanes is an application of the description of copoints in $G(\pm K_n)$ implied by Theorem II.2.1(h) (or see [37, Theorem 5.1(h)]). That the hyperplane in (b) is generated by its midpoints follows because a cutset of $-K_n$ is a balanced, connected, spanning subgraph of $\pm K_n$, hence has rank d. The parallelism follows from the farpoints contained in the hyperplane (and is metrically obvious given the spanning set of midpoints). □

Evidently, (c) and (d) are degenerate versions of (a) and (b) where the interesting action takes place within the opposite face and is governed there by (a) and (b). As a reminder of the natural isomorphism: the vertices are $v_1, \ldots, v_{d+1}$; $+e_{ij} \leftrightarrow$ farpoint on l_{ij}, $-e_{ij} \leftrightarrow$ midpoint on l_{ij}, and half edge $h_i \leftrightarrow$ vertex v_i.

The last statement generalizes the well known fact that, if we take one facet of s^d to be the base, the hyperplane generated by the midpoints of edges to the opposite vertex is parallel to the base.

We may look at (b) in a way that makes it more graphical. Take a spanning tree T in the 1-skeleton K_{d+1} and choose a sign $\sigma(e)$ for each edge. For each edge choose the midpoint if the edge is negative, the farpoint if it is positive. These points span a hyperplane $h(T, \sigma)$. This hyperplane meets every edge line of s^d in a midpoint or farpoint and we can say which.

Corollary 2.17. *An edge line l_{ij} meets $h(T, \sigma)$ in a midpoint or a farpoint according as the path in T from v_i to v_j is negative or positive.* □

Proof. Another application of Theorem 2.10, or of Corollary 2.16. □

In still another way of looking at Corollary 2.16(b) it is a variant and application of (a). The farpoints are cohyperplanar by definition. If we take a hyperplane h spanned by some midpoints and possibly farpoints, such as $h(T, \sigma)$, since no vertex is in h nothing prevents us from reaffinizing $\mathbb{P}^d(\mathbb{R})$ by throwing h to infinity. Reversing the process, we can think of h as h_∞ thrown by reaffinization to the ordinary hyperplane h. Let S be the set of edges corresponding to the generating midpoints and farpoints. Before reaffinization, when all these points were ideal, S had gain $+1$. Thus $\operatorname{clos} S$ was a subgraph $(K_{d+1}, +1)$, corresponding to the previous farpoints. After reaffinization S has gains ± 1, so the switching is by a ± 1-valued function; therefore $(K_{d+1}, +1)$ becomes a balanced gain graph (K_{d+1}, φ) where $\operatorname{Im} \varphi \subseteq \{\pm 1\}$. This means $\varphi^{-1}(-1)$ is a cutset of K_{d+1}. Therefore the midpoints and farpoints in h, which are described by φ, are as in (b). Thus we can deduce (b) from (a) by reaffinization, regarding any hyperplane spanned by midpoints and farpoints as a shifted ideal hyperplane.

2.6. **Ceva.**
Let
$$\Phi^* = (\Gamma, \varphi^*) \text{ where } \varphi^* = -1/\varphi.$$
and let $p^*(e) = p_{\Phi^*}(e)$ be the point associated to $e \in E(\Phi^*)$ under the definition (2.3). That is, $p^*(e) = v_{\mathbb{P}}$ if e is an unbalanced edge at v, and for a link
$$p^*(e) = \begin{cases} \frac{1}{1+\alpha}(\alpha v_{\mathbb{P}} + w_{\mathbb{P}}) & \text{if } \alpha = \varphi(e; v, w) \neq -1, \\ h_\infty \cap l_{vw} & \text{if } \alpha = -1. \end{cases} \tag{2.5}$$
Let $h_{\mathbb{P}}(e) = \operatorname{span}(N_{\mathbb{P}} \setminus \{v_{\mathbb{P}}\})$ if e is an unbalanced edge at v, and let
$$h_{\mathbb{P}}(e) = \operatorname{span}\big(N_{\mathbb{P}} \setminus \{v_{\mathbb{P}}, w_{\mathbb{P}}\} \cup \{p^*(e)\}\big)$$
if e is a link between v and w. Let $\mathcal{C}(\Phi) = \{h_{\mathbb{P}}(e) : e \in E\}$. This is the *Cevian representation* of $G(\Phi)$. We call $\varphi(e)$ the *Cevian coordinate* of $p^*(e)$. A hyperplane of the form $h_{\mathbb{P}}(e)$ is a *Cevian hyperplane* of $N_{\mathbb{P}}$; $p^*(e)$ is its *apex*. (Note that $h_{\mathbb{P}}(e)$ is the median hyperplane of edge $v_{\mathbb{P}}w_{\mathbb{P}}$ when $\varphi(e) = 1$ and is parallel to l_{vw} if $\varphi(e) = -1$.)

If (W, S) is a balanced subgraph of Φ^*, let $f : W \to F^*$ be a potential for $\Phi|(W, S)$ and let $f(W) = \sum_{w \in W} f(w)$. Call $\Phi|(W, S)$ *slim* if $f(W) = 0$. This property is independent of the choice of f; for example we may fix $v \in W$ and let $f(w) = \varphi(P_{vw})$ as in [39, p. 509], provided (W, S) is connected.

We call a subspace $\operatorname{span}(W_{\mathbb{P}})$, where $W \subseteq N$, a *facial subspace* for the basis $N_{\mathbb{P}}$. It is the projective subspace spanned by a face of the affine simplex whose vertex set is $N_{\mathbb{P}}$. If t is a flat of $\mathcal{C}(\Phi)$, let $A(t) = \{e \in E : h(e) \supseteq t\}$.

Theorem 2.18 (Generalized Ceva's Theorem)**.** *The set of flats of the projective family $\mathcal{C}(\Phi)$, ordered by reverse inclusion, is naturally isomorphic to* $\operatorname{Lat} G(\Phi)$.

A flat t lies in a facial subspace $\operatorname{span}(W_{\mathbb{P}})$ *if and only if $W \subseteq N_0(A(t))$. It lies in h_∞ if and only if every balanced component of $\Phi|A(t)$ is slim.*

Proof. The first statement is dual to the generalized Menelaus Theorem 2.10. If $X \subseteq \mathbb{A}^{N-1}$, let $X^\perp$ be the projective closure of the affine subspace $\{y \in \mathbb{A}^{N-1} : x \cdot y = 0 \text{ for all } x \in X\}$, where $x \cdot y = \sum_v x_v y_v$ if (x_v) and (y_v) are the affine coordinates of x and y in the basis $N_{\mathbb{P}}$. Then it is easily verified that $\{p(e)\}^\perp = h(e)$. The desired result follows.

The second statement is a consequence of the first. Let us enlarge Φ to $\Phi^\bullet$. Then $A^\bullet(\operatorname{span} W_{\mathbb{P}}) = E^\bullet{:}W$, which is contained in $A^\bullet(t)$ if and only if $W \subseteq N_0(A^\bullet(t)) = N_0(A(t))$.

The third part is established by the proof of [39, Theorem 13], which we reproduce here. We ask when every point x of F^N that projects into t lies in $h_0 : \sum x_v = 0$. The equations $x_w = \varphi(e; v, w)x_v$ for a link $e \in A(t)$ and $x_v = 0$ if an unbalanced edge exists in $A(t)$ at v force $x_v = 0$ for $v \in N_0(A(t))$ and $x_w = x_v\varphi(P_{vw})$ for v, w in a balanced component (W, S). Thus
$$\sum_{w \in W} x_w = x_v f(W).$$
If any $f(W) \neq 0$, we can find $x \notin h_0$. So all x must be in h_0 if, and only if, all $f(W) = 0$. □

Theorem 2.18 is an extension of the higher-dimensional Ceva theorem published in [10, 23]. It is also a generalization of the barycentric representation of contrabalanced Φ treated in [39, Section 3c].

In order to make our Cevian result completely geometrical we need to construct Φ from a geometric configuration. This we do as after Theorem 2.10 but with $\varphi^* = -1/\varphi$ replacing

φ there. Thus $\varphi(e; v, w) = \lambda^{-1} - 1 = -\beta/\alpha = d(p, w)/d(v, p)$. Then we can apply Theorem 2.10. (It was from the distance formula $\varphi = d(p, w)/d(v, p)$ that I originally derived Theorem 2.18. This formula still seems the easiest way to construct Cevian coordinates. I then found Theorem 2.10 by comparing the classical Menelaus and Ceva theorems. It may be worthwhile to look for other theorems of plane geometry in which something significant happens when a product around a triangle or polygon equals 1 (or -1), in the hope of finding new applications of gain graphs. I have tried this by making a quick survey of [1] but without noticing anything promising.)

If e and f are vw links, then from (2.4) the cross ratio of their Cevian points with respect to $v_{\mathbb{P}}$ and $w_{\mathbb{P}}$ is

$$(v_{\mathbb{P}}, w_{\mathbb{P}}; p^*(e), p^*(f)) = \frac{\varphi^*(e; v, w)}{\varphi^*(f; v, w)} = (v_{\mathbb{P}}, w_{\mathbb{P}}; p(e), p(f))^{-1}. \tag{2.6}$$

(Here we assume F is a field.) In particular, then, $p^*(e)$ and $p^*(f)$ are harmonic conjugates $\iff$ $p(e)$ and $p(f)$ are harmonically conjugate $\iff$ the gains sum to zero (in either Φ or Φ^*).

Example 2.11. Take the same 3-simplex as in Example 2.8 with the same points p_{1j}. The Cevian planes $h_{1j} = \operatorname{span}\{p_{1j}, v_k, v_l\}$ meet in a point q. The Cevian planes h_{ij} through q (for $j > i > 1$) have apices p_{ij}. We locate them by noting that the Cevian graph Φ determined by all the p_{ij} is balanced. The gains of the p_{1j} are $\varphi(e_{12}; v_1, v_2) = \varphi(e_{13}; v_1, v_3) = 2$, $\varphi(e_{14}; v_1, v_4) = 1/3$, so Φ is as in Figure 2.4. From this we locate p_{23} at the midpoint of edge v_1v_3; and p_{24} and p_{34} are one-seventh of the way from v_4 to v_2 and v_3, respectively.

By results of [36] and Theorem III.5.1, when $F = \mathbb{R}$ Theorem 2.10 implies that $\mathcal{M}(\Phi)$ has $\frac{1}{2}|\chi_\Phi(-1)|$ hyperplane separations. Dually, Theorem 2.18 implies a face enumeration formula for projective space.

Corollary 2.19. *Let $F = \mathbb{R}$ and let Φ be finite with at least one edge. The number of regions of $\mathcal{C}(\Phi)$ as a projective arrangement equals $\frac{1}{2}(-1)^n\chi_\Phi(-1)$. The face generating polynomial of $\mathcal{C}(\Phi)$ equals*

$$\frac{1}{2x}\left[(-1)^n w_\Phi(-x, -1) + x^{n-b(\Phi)}\right]. \qquad \square$$

The theorem also should imply a formula for faces, as well as one for bounded faces, applying to $\mathcal{C}(\Phi)$ restricted to affine space, but it is not yet possible to write them down, I believe because the true nature of slimness is not understood.

Problem 2.20. Axiomatize slimness.

One hopes for a graphically defined matroid $G_0(\Phi) = G(\Phi) \cup \{e_0\}$ that corresponds directly to the projective arrangement $\mathcal{C}_{\mathbb{P}}(\Phi) = \mathcal{C}(\Phi) \cup \{h_\infty\}$, in terms of whose characteristic and Whitney-number polynomials we would have simple formulas for the face numbers of $\mathcal{C}(\Phi)$ as an affine arrangment. For a start, clearly $\{S \in \operatorname{Lat}\Phi : \text{ every balanced component of } (N, S) \text{ is slim}\}$ is a modular filter. Which modular filters are produced in this way?

Example 2.12 (*Medians and parallels*). Let $d \geq 2$ and take a d-simplex s^d in $\mathbb{A}^d(\mathbb{R})$ with vertices $N = \{v_1, \ldots, v_d, v_{d+1}\}$. The *median hyperplane* m_{ij} is the Cevian hyperplane of the midpoint of the segment v_iv_j, that is, $\operatorname{span}\big(N \setminus \{v_i, v_j\} \cup \text{midpoint}\big)$. The *Cevian parallel* or *paracevian* c_{ij} is the Cevian of the farpoint on l_{ij}, that is, it contains $N \setminus \{v_i, v_j\}$ and is

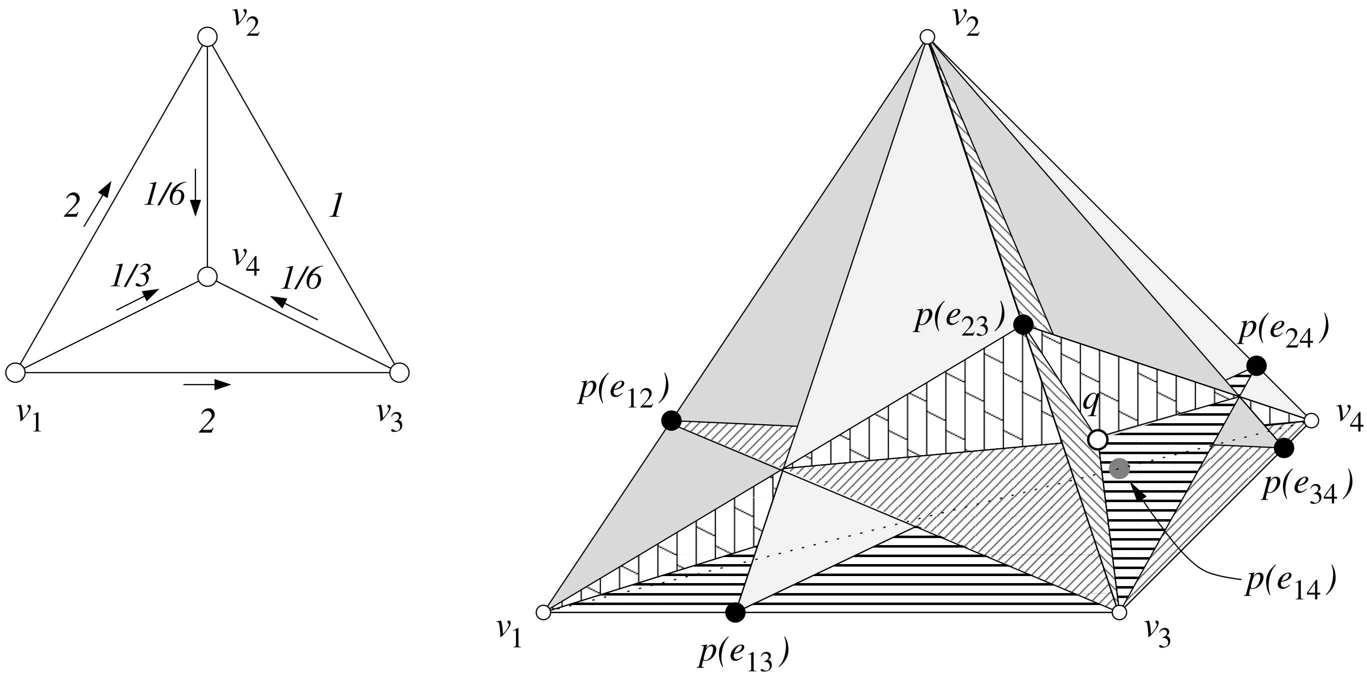

FIGURE 2.4. A real multiplicative gain graph Φ and the corresponding Cevian apices and planes (concurrent at q) of a tetrahedron in $\mathbb{A}^3(\mathbb{R})$, as in Example 2.11.

parallel to l_{ij}. We can describe the geometry of the arrangement $\mathcal{P}$ of all paracevians, that of $\mathcal{P}^\bullet$, the arrangement of all paracevian and facet hyperplanes, that of $\mathcal{M}$, the arrangement of all median hyperplanes, and that of their combination $\mathcal{P}^\bullet \cup \mathcal{M}$.

Harary proved that a balanced signed graph Σ has a bipartition $\{X, Y\}$ of its nodes such that an edge is negative if and only if it has one end in X and the other in Y [19]. (X or Y may be void.) This is a *Harary bipartition* of Σ. It is unique if Σ is connected. We say Σ is *evenly bipartitioned* if $\#X = \#Y$.

Corollary 2.21. *The intersection lattice $\mathcal{L}(\mathcal{P}^\bullet \cup \mathcal{M})$ is $\operatorname{Lat} G(\pm K^\bullet_{d+1})$, with half edges corresponding to facet hyperplanes, positive edges to median hyperplanes, and negative edges to paracevians.*

An intersection flat is contained in the ideal hyperplane if and only if it corresponds to an edge set in which every balanced component is evenly bipartitioned.

Proof. The first part is a special case of Theorem 2.18. For the second we note that a potential for a balanced signed graph Σ with Harary bipartition $\{X, Y\}$ is $f\big|_X \equiv +1$, $f\big|_Y \equiv -1$. $\square$

The chromatic polynomials of some interesting subgraphs of $\pm K_n^\bullet$ are known. Mainly from [38, Equations (1.1), (3.1), and (5.7)], we have:

$$\chi_{+K_n}(\lambda) = (\lambda)_n, \qquad \chi_{+K_n^\bullet}(\lambda) = (\lambda-1)_n,$$

$$\chi_{-K_n}(\lambda) = \sum_{i=1}^{n} \big[S(n,i) + nS(n-1,i)\big] 2^i \left(\frac{\lambda-1}{2}\right)_i,$$

$$\chi_{-K_n^\bullet}(\lambda) = \sum_{i=1}^{n} S(n,i) 2^i \left(\frac{\lambda-1}{2}\right)_i,$$

$$\chi_{\pm K_n}(\lambda) = 2^{n-1} \left(\frac{\lambda}{2}\right)_{n-1} \big[\lambda - n + 1\big],$$

$$\chi_{\pm K_n^\bullet}(\lambda) = 2^n \left(\frac{\lambda-1}{2}\right)_n.$$

Here $(\lambda)_n$ and $S(n,i)$ are as in Example 2.3. Therefore by Corollary 2.19 we can write down $r_{\mathbb{P}}$, the number of regions of $\mathcal{P}^\bullet \cup \mathcal{M}$ and various subarrangements, regarded as *projective* arrangements. (We assume $d \geq 1$ so Cevians exist.) The numbers are

$$r_{\mathbb{P}}(\mathcal{M}) = \frac{1}{2}(d+1)!, \qquad r_{\mathbb{P}}(\mathcal{M}^\bullet) = \frac{1}{2}(d+2)!,$$

$$r_{\mathbb{P}}(\mathcal{P}) = \sum_{i=1}^{d+1} \big[S(d+1,i) + (d+1)S(d,i)\big] 2^{i-1} i! (-1)^{d+1-i},$$

$$r_{\mathbb{P}}(\mathcal{P}^\bullet) = \sum_{i=1}^{d+1} S(d+1,i) 2^{i-1} i! (-1)^{d+1-i},$$

$$r_{\mathbb{P}}(\mathcal{P} \cup \mathcal{M}) = 2^{d-1}(d+1)!, \qquad r_{\mathbb{P}}(\mathcal{P}^\bullet \cup \mathcal{M}) = 2^d (d+1)!.$$

(The numbers for the all-negative complete graphs are similar to the numbers in Example 2.3, for the very good reason that the gain graphs are the same.) The face numbers in projective space can be obtained from formulas for Whitney-number polynomials in [38]: e.g., see [38, (5.9)] (with [38, (1.2)]) for $-K_n$ and $-K_n^\bullet$. I omit them because they are complicated, but at least they are known.

The same cannot be said for the number of regions and faces in affine space. We cannot apply formulas because we do not know enough. In principle, the numbers could be found either by adding h_∞ to the arrangement in projective space and applying [36, Theorem B] (this is the approach suggested at Problem 2.20), or by deleting all ideal flats from $\mathcal{L}(\mathcal{C}(\Phi))$, leaving the semilattice of an affine arrangement to which one can apply [36, Theorem C]. The latter only requires identifying the ideal flats, which are characterized in Corollary 2.21, and removing them from the intersection lattice, leaving a lower ideal (a geometric semilattice, in fact) whose characteristic polynomial is perhaps knowable.

Problem 2.22. Find the numbers of regions of $\mathcal{P}$, $\mathcal{P}^\bullet$, $\mathcal{P}^\bullet \cup \mathcal{M}$, etc., as affine arrangements; also, the number of k-faces for every k.

One cannot escape noticing that the intersection lattice of $\mathcal{P}^\bullet \cup \mathcal{M}$, Lat $G(\pm K_n^\bullet)$, is identical to that of the root system arrangement B_n^*, which is merely the real canonical hyperplanar bias representation $\mathcal{H}(\pm K_n^\bullet)$. In fact, $\mathcal{P}^\bullet \cup \mathcal{M}$ as a projective arrangement is isomorphic to $\mathcal{H}(\pm K_n^\bullet)$ projected into $\mathbb{P}^{n-1}(\mathbb{R})$.

The Menelæan representation leads to coordinates on the system of projective lines spanned by a basis. The Cevian representation leads to a coordinate system for the whole set of projective points. Take $\Phi = F^* K_N^\bullet$: then each point in $\mathbb{P}^{N-1}$ outside the hyperplanes spanned by $N_\mathbb{P}$ corresponds to a unique balanced copoint of $G(\Phi)$. The other projective points correspond to the unbalanced copoints in $G(\Phi)$. Thus in principle a part of projective geometry should be expressible through biased graphs. This line of thought, though a variant of the usual theory of coordinatization by geometric nets, seems to have its own potential. I hope to develop it elsewhere.

3. Abstract logic and geometry of the bias matroid

We can represent the bias matroid of any gain graph, regardless of what its gain group may be, by abstracting from the hyperplane geometry of Corollary 2.2 the concept of formal equations with at most two terms, briefly *two-term equations*. These are equations of the forms $x_v = \alpha x_w$, $x_v = 0$, and (for completeness) $\beta x_v = \alpha x_w$, in variables x_v indexed by a set N and with constants α, β, and 0 where α and β belong to a group $\mathfrak{G}$ and $0 \notin \mathfrak{G}$. The rules of implication amongst these equations are the usual rules of equality and, in addition,

(i) $1x_v = x_v$, where 1 is the group identity,
(ii) $x_v = \alpha x_w \iff \beta x_v = (\beta\alpha) x_w$ for $\alpha, \beta \in \mathfrak{G}$,
(iii) $x_v = \alpha x_v$ for $\alpha \neq 1 \iff x_v = 0$.

Let E_N be the set of all equations of the three types that involve variables x_v for $v \in N$. Suppose $E \subseteq E_N$. We call a subset $S \subseteq E$ *logically closed in* E if no equation in $E \setminus S$ can be deduced from S. (It is permissible to use equations not in E at intermediate steps in a deduction.)

Suppose we have a set $E \subseteq E_N$. We may represent it by a gain graph. We take N for the node set, $\mathfrak{G}$ for the gain group, and E for the edge set. An equation $x_v = 0$ is represented by a half edge at v. An equation $\beta x_v = \alpha x_w$ (where v may equal w) is represented by a link or loop between v and w whose gain is $\varphi(e; v, w) = \beta^{-1}\alpha$. Let $\Lambda(N, E)$ denote this gain graph. Note that $\Lambda(N, E_N)$ is $\mathfrak{G} K_N^\bullet$ with multiple edges due to the several equivalent equations $\beta x_v = \beta\alpha x_w$, $\beta \in \mathfrak{G}$; we shall for simplicity treat these multiple edges as the same.

Theorem 3.1. *Let E be a set of two-term equations in variables x_v, $v \in N$. A subset $S \subseteq E$ is logically closed in E precisely when S is closed in $G(\Lambda(N, E))$.*

Proof. Let $\operatorname{clos}_\Lambda(S)$ denote the logical closure in E. It is obvious that no deduction requires any variables besides the x_v for $v \in N$.

We show first that $\operatorname{clos}_\Lambda(S) \supseteq \operatorname{clos}_G(S)$. In an unbalanced component of S we can obviously deduce $x_v = 0$ for all nodes. Let e be a link between v and w, where we have already deduced $x_v = 0$ and $x_w = 0$. Then $x_w = \alpha x_w$ for all α, so $x_v = 0 = x_w = \alpha x_w$ implies $x_v = \alpha x_w$. Thus $\operatorname{clos}_\Lambda(S) \supseteq E{:}N_0(S)$.

In a balanced component of S we have

$$\operatorname{clos}_G(S) = \operatorname{bcl}(S) = S \cup \{e \in E : e \in C \subseteq S \cup \{e\} \text{ for some balanced circle } C\}.$$

Let e in this statement be a link from v_0 to v_k and let $C \setminus e \subseteq S$ be the path $e_1 e_2 \cdots e_k$ where e_i links v_{i-1} to v_i. We deduce $x_{v_0} = \varphi(e_1 \cdots e_i) x_{v_i}$ successively for $i = 1, 2, \ldots, k$, arriving at $x_{v_0} = \varphi(e) x_{v_k}$, which is the equation corresponding to e. If e was a balanced loop at v, the corresponding equation $x_v = 1x_v$ follows from (i).

Now we show that $\operatorname{clos}_G(S) \supseteq \operatorname{clos}_\Lambda(S)$. A deduction consists of a sequence of applications of Rules (i)–(iii) and the three laws of equivalence relations. We reinterpret each rule graphically in $\mathfrak{G}K_N^\bullet$. Rule (i) corresponds to the statement that a balanced loop is in any G-closed set, which is true. Rule (iii) translates to say that, if one unbalanced edge at v is in the G-closure, then all are; this is correct. The symmetric law, with Rule (ii), says that a link or loop should have the inverse gain if its direction is reversed. The transitive law and Rule (ii) say that, if the closure in $\mathfrak{G}K_N^\bullet$ contains edges e from v to w and f from w to x, then it contains an edge g from v to x with gain $\varphi(g) = \varphi(e)\varphi(f)$. This is true of G-closure in $\mathfrak{G}K_N^\bullet$. Thus, the logical closure L of S within E_N lies in the bias closure T of S within $\mathfrak{G}K_N^\bullet$. Since $\operatorname{clos}_\Lambda(S) = E \cap L$ and $\operatorname{clos}_G(S) = E \cap T$, we have the desired inclusion. $\square$

A more concrete version of this idea, closer to the geometry in Corollary 2.2, was suggested to me by Jay Sulzberger [32]. It is a kind of *permutation gain graph*, where the gain group is a permutation group. Given Φ, let Z be a set, containing a special element $\hat{0}$, on which $\mathfrak{G}$ acts as a permutation group so that each group element leaves fixed only $\hat{0}$. (For example F may be a skew field, $Z = F$, and $\mathfrak{G} \subseteq F^*$.) Let

$$h(e) = \begin{cases} \{x \in Z^N : x_v = \varphi(e; v, w)x_w\} & \text{if } \nu_\Gamma(e) = \{v, w\}, \\ \{x \in Z^N : x_v = \hat{0}\} & \text{if } e \text{ is a half edge at } v, \\ Z^N & \text{if } e \text{ is a loose edge or balanced loop,} \end{cases}$$

and let $X(S) = \bigcap\{h(e) : e \in S\}$ if $S \subseteq E$. If $Y \subseteq Z^N$, let $E(Y) = \{e \in E : h(e) \supseteq Y\}$.

Theorem 3.2. *The set $\{X(S) : S \subseteq E\}$ ordered by reverse inclusion is isomorphic to* $\operatorname{Lat} G(\Phi)$. *We have* $\operatorname{clos}_G S = E(X(S))$. $\square$

4. Geometry of the lift matroid

A *lift representation* of Ω is a vector, affine, or projective representation of $L(\Omega)$. Amongst all lift representations, one kind is canonical.

4.1. Canonical representations.

Suppose the gain group $\mathfrak{G}$ of Φ is an additive subgroup of a skew field F (written $\mathfrak{G} \leq F^+$, F^+ being the additive group of F). For each $e \in E_0$ we define a vector $z_\Phi(e)$, or simply $z(e)$, $\in F^{1+N} = \{(x_0, x) : x_0 \in F,\ x \in F^N\}$ by

$$z_\Phi(e) = \begin{cases} (\varphi(e; v, w), \hat{w} - \hat{v}) & \text{if } \nu_\Gamma(e) = \{v, w\}, \\ (1, 0) & \text{if } e \text{ is a half edge or the extra point } e_0, \\ (0, 0) & \text{if } e \text{ is a loose edge or balanced loop,} \end{cases}$$

where $\hat{v}$ denotes the unit basis vector in the v direction. The vector $z(e)$ is well defined up to negation; that is enough for our theorem. We call any z_Φ a *standard lift representation of* Φ.

Take an F-vector space Λ that contains F^{1+N} or is extendible to contain it. We call any function $f : E \to \Lambda$ such that each $f(e)$ is a non-zero scalar multiple of $z(e)$ a *canonical lift representation* of Φ. If $f : E_0 \to \Lambda$ is a function with the same property for each $e \in E_0$, f is a *canonical complete-lift representation* of Φ. (In full, we say, e.g., *canonical linear* (or *vector) lift*, or *complete-lift, representation over F.)* These terms are justified by the first part of Theorem 4.1.[6]

For the second part, first we need to know the effect on z_Φ of switching Φ by η. It is

$$z_{\Phi^\eta} = \left[\begin{array}{c|c} 1 & \eta^{\mathrm{T}} \\ \hline 0 & I_N \end{array}\right] z_\Phi, \tag{4.1}$$

where we regard z and η as column vectors ($z \in F^{1+N}$ and $\eta \in F^N$) and I_N is the $N \times N$ identity matrix. This follows from the expressions $z_{\Phi^\eta}(e) = \big(-\eta(v)+\varphi(e;v,w)+\eta(w),\ \hat{w}-\hat{v}\big)$ if e has endpoints v and w; and $z_{\Phi^\eta}(e) = (1,0)$ for a half edge, $(0,0)$ for a loose edge. We conclude that the canonical lift representation of a switching class $[\Phi]$ is well defined up to projective equivalence. (But not conversely; see Section 4.3.) A *canonical [complete] lift representation* of a switching class $[\Phi]$ is any canonical [complete] lift representation of any switching Φ^η.

Now we need to define the restriction $f|_S$ and contraction f/S of a canonical representation of $L_0(\Phi)$ (or $L(\Phi)$). The restriction is elementary: merely restrict f to $S \cup \{e_0\}$ (or S). The definition of the contraction when S is balanced is based on the fact that $f(e) = \alpha_e z_\Phi(e)$ where α_e is an arbitrary nonzero scalar. f/S will be a function $E_0 \setminus S \to F^{1+\pi_{\mathrm{b}}(S)}$. Choose a switching function η for which $\varphi^\eta|_S \equiv 0$. Write ξ_0 for the leading coordinate of $\xi \in F^{1+N}$ or $F^{1+\pi_{\mathrm{b}}(S)}$. The definition of f/S is

$$\begin{aligned}(f/S)(e)_V &= \sum_{u\in V} f(e)_u = \alpha_e \sum_{u\in V} z_\Phi(e)_u \quad \text{if } V \in \pi_{\mathrm{b}}(S),\\ (f/S)(e)_0 &= \alpha_e z_{\Phi^\eta}(e)_0.\end{aligned}$$

Thus $(f/S)(e) = \alpha_e(z_{\Phi^\eta}/S)(e)$. By our definition, f/S is well defined on $[\Phi]$ up to projective equivalence, though its exact value depends on the choice of η. Furthermore, it is clear that

$$(z_{\Phi^\eta})/S = z_{\Phi^\eta/S}.$$

If S is unbalanced, then f/S is constructed by deleting the 0th coordinate and contracting S as with the canonical bias representation of the identity-gain graph $(\Gamma, 1)$.

Theorem 4.1. *Let $\Phi = (\Gamma, \varphi, \mathfrak{G})$ be a gain graph and F a skew field containing $\mathfrak{G}$ as an additive subgroup.*

(a) *The linear dependence matroid of the vectors $z_\Phi(e)$ and $z_\Phi(e_0) \in F^{1+N}$ is isomorphic under the mapping z_Φ to the complete lift matroid $L_0(\Phi)$.*

(b) *For any canonical lift [or, complete lift] representation f and any $S \subseteq E$, $f\big|_S$ is a canonical lift [or, complete lift] representation of $\Phi|S$ and, if S is balanced, f/S is a canonical lift [or, complete lift] representation of $[\Phi/S]$.*

Proof of (a). We treat a half edge as an unbalanced loop. If we suppress the first coordinate, what remains is the standard representation of the graphic matroid $G(\Gamma)$ (neglecting the extra point). Thus a linearly dependent set must contain a circle or the extra point. A set $\{z(e) : e \in S\}$, where $S \subseteq E$ contains a unique circle C, is linearly dependent if and only if it is balanced, for in F^N the only linear dependence among the elements of S is the sum of the edges in C (suitably oriented), which in F^{1+N} has first coordinate $\sum_C \varphi(e)$. If C is

[6]Our construction is a more precise version of a standard way to construct geometrical representations of lifts (or "coextensions") of a matroid. The best reference I can find is Figure 7.6 and the second part of Proposition 7.4.17(3) in [8].

unbalanced, the $z(e)$ for $e \in S$ span the vector $(1,0) = (\sum_C z(e))/(\sum_C \varphi(e))$. Therefore, a set $T \subseteq E \cup \{e_0\}$ is dependent if it contains two unbalanced figures, or one such figure and e_0. But if it is a forest together with either e_0 or one more edge forming an unbalanced figure, it is obviously independent. This completes the proof. $\square$

Proof of (b). The restriction is trivial. For contraction of a balanced set S, since $(f/S)(e) = \alpha_e(z_{\Phi^\eta}/S)(e)$ and $(z_{\Phi^\eta})/S = z_{\Phi^\eta/S}$, f/S is a canonical lift representation of $[\Phi/S]$. If S is unbalanced then by Theorem 2.1(b) f/S represents $G((\Gamma,1)/S) = G(\Gamma/S)$. $\square$

I do not know whether Theorem 4.1(a) has previously appeared as such. But the matrix of the representation in the real case (without e_0) is well known since it arises in network flow problems with one linear side constraint; cf. [20]. The binary representation matrix (with e_0) of a signed graph was used by Gerards in important work on signed graphs [16].

Example 4.1. In Figure 4.1(a) is a gain graph of order $n = 3$ with gains in $\mathbb{Z}$, the additive group of integers. The balanced circles are $C_1 = \{3e_{12}, 0e_{23}, 3e_{13}\}$, $C_2 = \{1e_{12}, 2e_{23}, 3e_{13}\}$, and $C_3 = \{0e_{12}, 2e_{23}, 2e_{13}\}$, since, e.g., $\varphi(C_1) = 3+0+(-3) = 0$. So, $\langle\Phi\rangle = (\|\Phi\|, \{C_1, C_2, C_3\})$. (The notation for edges is as in Example 2.1 but, of course, additive. Thus for instance $2e_{13} = (-2)e_{31}$.)

The canonical representation vectors for $L(\Phi)$ (and $L_0(\Phi)$) in $\mathbb{R}^{1+3}$ are

$$z(h_1) = z(e_0) = b_0 = (1;0,0,0,), \quad z((-1)e_{22}) = (-1)b_0 + b_2 - b_2 = (-1;0,0,0),$$
$$z(0e_{12}) = (0;-1,1,0), \quad z(1e_{12}) = (1;-1,1,0), \quad z(3e_{12}) = (3;-1,1,0),$$
$$z(0e_{23}) = (0;0,-1,1), \quad z(2e_{23}) = (2;0,-1,1),$$
$$z(2e_{13}) = (2;-1,0,1), \quad z(3e_{13}) = (3;-1,0,1);$$

they are shown in Figure 4.1(b). The corresponding canonical hyperplanar lift representation (see Corollary 4.5) is four-dimensional, which is awkward to display; nevertheless, because $(0;1,1,1)$ lies in every hyperplane one could represent it by its faithful three-dimensional cross section in $\mathbb{R} \times s$ where $s = \{x \in \mathbb{R}^3 : x_1 + x_2 + x_3 = 0\}$. (The representing vectors lie in $\mathbb{R} \times s$ and they span it because $\operatorname{rk} L(\Phi) = 3$.) However, we eschew this in favor of depicting the affinographic representation (Corollary 4.5).

To illustrate the construction we treat $3e_{12}$. The dual hyperplane of $z(3e_{12}) = (3;-1,1,0)$ is $x_2 - x_1 = 3x_0$ in $\mathbb{R}^{1+3}$. (For duality see before Corollary 4.5.) Setting $x_0 = 1$ gives the affinographic hyperplane $x_2 - x_1 = 3$ in the affine space $\{1\} \times \mathbb{R}^3$. (An edge like h_1 or $(-1)e_{22}$ requires special consideration: $z(h_1) = (1;0,0,0)$ dualizes to $0 = 1x_0$, and setting $x_0 = 1$ gives an inconsistency. This means that the affinographic hyperplane is the ideal hyperplane h_∞ in the projective completion of $\{1\} \times \mathbb{R}^3$.) Thus, the affinographic arrangement $\mathcal{A}(\Phi)$ consists of 2-dimensional hyperplanes; with h_∞ we have its projectivization $\mathcal{A}_{\mathbb{P}}(\Phi)$. But these hyperplanes can be reduced further, for all contain the point $(1;1,1,1)$. Cross-sectioning by s gives the affine planar arrangement $\mathcal{A}^s$ shown in Figure 4.1(c). $\mathcal{A}(\Phi)$ itself is obtained by coning on a point outside the plane of the figure. (That is, each line is raised to a plane by joining it with the outside point.) With the addition of the ideal line l_∞, the diagram shows $\mathcal{A}_{\mathbb{P}}(\Phi)^s$, that is, $\mathcal{A}_{\mathbb{P}}(\Phi)$ sectioned by the plane of the figure. Again, one obtains $\mathcal{A}_{\mathbb{P}}(\Phi)$ by coning on a point outside the plane.

The gain group $\mathbb{Z}$ has many embeddings in $\mathbb{R}^+$, the additive group of real numbers, but all are related by scaling the 0^{th} coordinate. Thus any projectively inequivalent representations of $L(\Phi)$ are either canonical lift representations of other gain graphs Φ' for which $\langle\Phi'\rangle = \langle\Phi\rangle$,

(a)

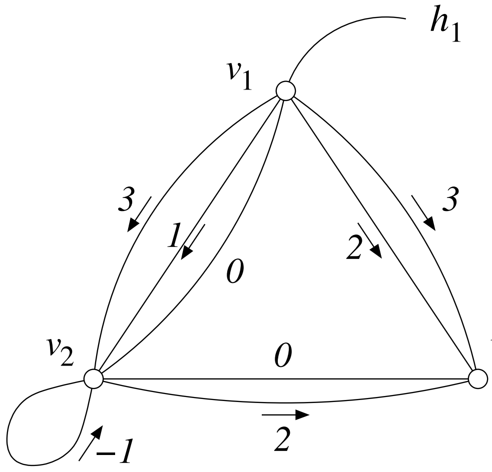

(b)

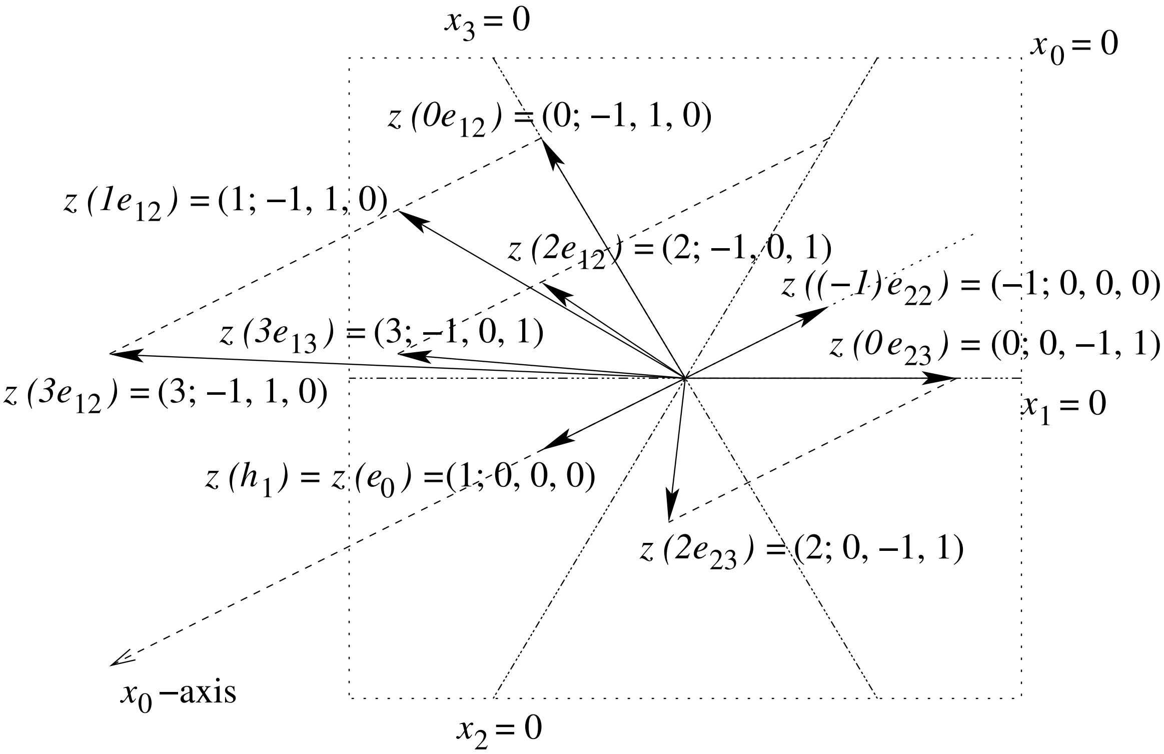

FIGURE 4.1. A $\mathbb{Z}$-gain graph Φ (a), a canonical real vector lift representation (b), and (c) a generic planar cross section of the affinographic representation (without l_∞) or its projectivization (with l_∞).

(c)

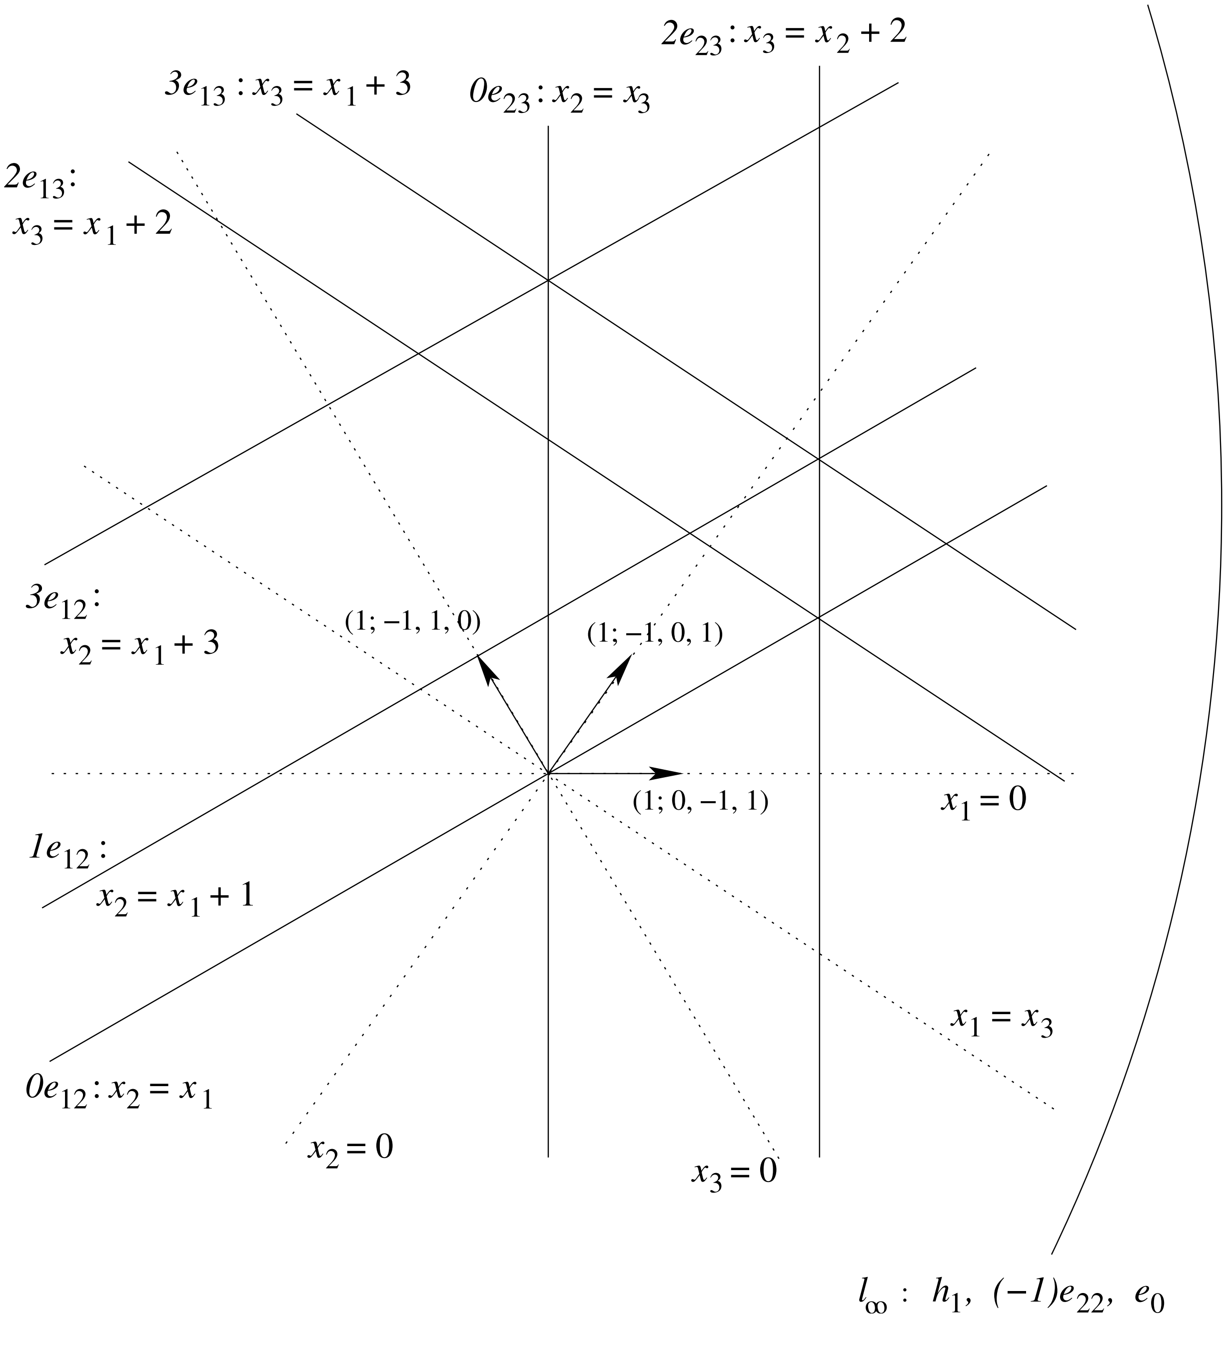

of which there are many (by a discussion similar to that in Example 2.1), or are noncanonical. In this example, though, where Φ has an unbalanced edge e, so that $e_0 \in \operatorname{clos}_L e$ and $L(\Phi)$ is essentially $L_0(\Phi)$, all representations are canonical lift representations by Proposition 4.3. (One can also deduce that conclusion from Theorem 7.1.)

Example 4.2 (*Spikes; continuation of Example 2.2*). We define a *spike* as a matroid $L(2C_n, \mathcal{B})$ where $n \geq 3$ and $\mathcal{B}$ is a Hamiltonian bias; $L_0(2C_n, \mathcal{B})$ is a *complete spike.* A *free* *spike* is a bicircular lift matroid $L(2C_n, \varnothing)$ with $n \geq 3$. Spikes, like swirls, are crucial examples in matroid representability (complete free spikes are the matroids N_r of [30, Section 5], and [15, Section 7] employs two kinds of spikes, free ones and those with two complementary circuit hyperplanes), but it is apparently a new observation that spikes are lift matroids that arise from the same biased graphs as swirls, whence their properties are analogous.

We may visualize a spike as a matroid consisting of n 2-point lines, concurrent in a base point p_0 not part of the spike (but it is part of the complete spike!) but otherwise independent; that is, the spike has rank n. There may in addition be circuit hyperplanes involving one point from each 2-point line. This is the definition in [15]; we see immediately that a spike is a lift matroid as we described.

A free spike, arising from a contrabalanced graph, is representable over any sufficiently large field. Part VII has crude bounds on how large.

By Theorem 7.1 a spike $L(2C_n, \mathcal{B})$ with $n \geq 4$ is representable over F if and only if $(2C_n, \mathcal{B})$ has gains in F^+. If $(2C_n, \mathcal{B})$ has no gains at all, then its spike *and* its swirl are not representable in any vector space over any field. By the same theorem any representation of a spike is a canonical lift representation. (It is easy to convert the representation in [30, p. 340] to canonical form by adding a new row and performing simple row operations.) The unique binary spike is $L(\pm C_n)$, which is the column-dependence matroid of the matrix $(I \mid J - I)$, where J is the square all-ones matrix (see [12a], [12b, p. 245]). Its uniqueness is obvious from the fact that the only sign-biased spike is $\langle \pm C_n \rangle$.

Example 4.3 (*Balanced*). If Φ is balanced, as in Example 2.5, then $L(\Phi)$ is uniquely representable so a canonical lift representation is any representation in a vector or projective space.

Corollary 4.2. *If Φ is finite and has no half edges, and if $F = \mathbb{R}$, then the set $Z = \{z(e) : e \in E\}$ has $|\chi^{\mathrm{b}}_\Phi(-1) - \frac{1}{2}\chi_\Gamma(-1)|$ hyperplane separations. The set $Z \cup \{z(e_0)\}$ has $|\chi^{\mathrm{b}}_\Phi(-1)|$ separations.*

Proof. We rely on Theorem III.5.2 together with results on counting separations by hyperplanes cited in Section 1. □

One cannot expect a complete converse to Theorem 4.1(a), saying that any representation of $L(\Phi)$ is canonical. For instance, $L(\pm K_3) = G(\pm K_3) \cong G(K_4)$, which is linearly representable over every field; but $\langle \pm K_3 \rangle$ has a canonical lift representation only in characteristic 2. (This is the same example we used for Theorem 2.1's converse. See Example 2.4.) Nonetheless there is a converse for $L_0(\Phi)$ and (which is equivalent) for $L(\Phi)$ if Φ contains an unbalanced edge.

Proposition 4.3. *Let $\Omega = (\Gamma, \mathcal{B})$ be a biased graph and $f : E \cup \{e_0\} \to \Lambda$ a vector representation of $L_0(\Omega)$ over a skew field F. Then there is a gain graph $\Phi = (\Gamma, \varphi, F^+)$ such that $\langle \Phi \rangle = \Omega$ and f is a canonical complete lift representation of Φ.*

Proof. Again we treat a half edge like an unbalanced loop. We assume Λ is large enough to contain F^{1+N}.

Let $\hat{v}_0 = f(e_0)$ and let $f' = f \circ g$, where $g : \Lambda \to \Lambda_0 = \Lambda / \operatorname{span}(\hat{v}_0)$ is the projection. Then f' is a representation of $G(\Gamma)$. Since such a representation is projectively unique, we can find a coordinate system for Λ_0 in which $\hat{N} = \{\hat{v} : v \in N\}$ is part of a basis and $f'(e) = \alpha(e)(\hat{w} - \hat{v})$ if $\nu_\Gamma(e) = \{v, w\}$.

Now express Λ as the direct sum $\operatorname{span}(\hat{v}_0) \oplus \Lambda_0$ in some way and let $\varphi(e; v, w)$ be the value α such that $f(e) - \alpha\hat{v}_0 \in \Lambda_0$. That defines Φ. Since $f(e) = (\alpha; \hat{w} - \hat{v})$, we have a canonical complete lift representation of Φ in the subspace $\operatorname{span}(\hat{N} \cup \hat{v}_0)$. Since f is a canonical complete lift representation of Φ and a representation of $L_0(\Omega)$, the identity mapping of edges is an isomorphism of the matroids. Since the underlying graphs are also the same, Φ and Ω must have the same balanced circles. $\square$

Due to Proposition 4.3 we can define a (linear) *canonical complete lift representation of* Ω as either a canonical complete lift representation of any gain graph Φ for which $\langle\Phi\rangle = \Omega$, or equivalently as any linear representation at all of $L_0(\Omega)$. A *canonical lift representation of* Ω is either a canonical lift representation of any gain graph Φ whose bias is Ω, or equivalently the restriction to E of any representation of $L_0(\Omega)$. Based on this equivalence we can give coordinate-free projective and affine definitions. A *canonical projective* [or, *affine*] *complete lift representation* of Ω is a representation of $L_0(\Omega)$ in any projective [affine] space, including noncoordinatizable projective or affine planes, and a *canonical projective* [*affine*] *lift representation* is the restriction to E of a projective [affine] representation of $L_0(\Omega)$.

Problem 4.4 (*Fundamental representation question*). (a) Characterize the biased graphs that have a noncanonical lift representation (not counting canonical bias representations when $L(\Omega) = G(\Omega)$); most especially, one that is not binary. (b) What kinds of noncanonical representation can exist in these cases?

Example 4.4 (*Too thick; continuation of Example 2.7*). A biased graph Ω as in Example 2.7 has no canonical lift representation over $\mathbb{F}_{k-1}$ (assuming $k-1 = q$, a prime power) because $L_0(\Omega)$ has a $k+1$-point line and $k+1 > q+1$. Nonetheless, $L(\Omega)$ might be representable over $\mathbb{F}_{k-1}$: one such example is $L(\Omega_k)$. As in Example 2.7, this noncanonicality is reparable by extending the field.

There is also, of course, a hyperplane representation dualizing Theorem 4.1; indeed, there are several. First, one can dualize the canonical [complete] lift representation in F^{1+N}, so that e corresponds to $z^*_\Phi(e) =$ the hyperplane $x_w - x_v = \varphi(e; v, w)x_0$ and e_0 or a half edge to $z^*_\Phi(e_0) =$ the hyperplane $x_0 = 0$; that gives a linear arrangement $\bar{A}(\Phi)$. This is the *canonical linear* (or *homogeneous*) *hyperplanar* [*complete*] *lift representation* of Φ. (Note that this duality is abnormal because b_0 dualizes to $-x_0$. The ordinary dual hyperplane would be $x_w - x_v + \varphi(e; v, w)x_0 = 0$. Our duality is chosen for the sake of nice equations in the affinographic representation.) Then, one can treat the equations as involving homogeneous coordinates in the projective space $\mathbb{P}^N(F)$, with infinite hyperplane $h_\infty : x_0 = 0$ corresponding to e_0; thus we have a projective arrangement $\mathcal{A}_{\mathbb{P}}(\Phi)$ (which contains h_∞). Finally, by removing that hyperplane we have a representation by an affine hyperplane arrangement $\mathcal{A}(\Phi)$ in which the hyperplane corresponding to e has equation $x_w - x_v = \varphi(e; v, w)$. This is the *affinographic representation* of Φ (for which see also the end of Section 4.5) and $\mathcal{A}_{\mathbb{P}}(\Phi)$ is the *projectivized affinographic representation*. Real affinographic representations of $\mathbb{Z}$-gain graphs have lately become popular; see Example 4.5.

Corollary 4.5. (a) *Under the mutually inverse correspondences*

$$A \in \operatorname{Lat} L_0(\Phi) \mapsto \bigcap_{e \in A} h(e) \quad \textit{and} \quad t \in \operatorname{Lat} L_0(\Phi) \mapsto \{e \in E_0 : h(e) \supseteq t\}$$

(where $h(e)$ denotes the hyperplane representing e in $\bar{\mathcal{A}}(\Phi)$, $\mathcal{A}_{\mathbb{P}}(\Phi)$, or $\mathcal{A}(\Phi)$ as appropriate, for each $e \in E_0$ or just $e \in E$ in the affine case), the intersection lattices $\mathcal{L}(\bar{\mathcal{A}}(\Phi))$ and $\mathcal{L}(\mathcal{A}_{\mathbb{P}}(\Phi))$ are isomorphic to $\operatorname{Lat} L_0(\Phi)$ *and the intersection semilattice $\mathcal{L}(\mathcal{A}(\Phi))$ is isomorphic to* $\operatorname{Lat^b}\Phi$.

(b) *Suppose Φ is finite and $F = \mathbb{R}$. The numbers of regions, if Φ is unbalanced, are: $2|\chi^{\mathrm{b}}_{\Phi}(-1)|$ of $\bar{\mathcal{A}}(\Phi)$, and $|\chi^{\mathrm{b}}_{\Phi}(-1)|$ of $\mathcal{A}(\Phi)$ and $\mathcal{A}_{\mathbb{P}}(\Phi)$.*

(c) *Suppose Φ is finite and $F = \mathbb{C}$. The Poincaré polynomial of the complement $\mathbb{C}^{n+1} \setminus \bigcup \bar{\mathcal{A}}(\Phi)$ is equal to $(y+1)(-y)^n\chi^{\mathrm{b}}_{\Phi}(-1/y)$. That of the complement $\mathbb{A}^n(\mathbb{C}) \setminus \bigcup \mathcal{A}(\Phi)$ equals $(-y)^n\chi^{\mathrm{b}}_{\Phi}(-1/y)$.*

Proof. (a) Mainly a straightforward dualization of Theorem 4.1. Note that $\operatorname{Lat^b}\Phi$ corresponds to the affine intersection flats of the projective arrangement. Thus the projective arrangement representing $\operatorname{Lat} L_0(\Phi)$ is, as the notation suggests, the projectivization of the affine arrangement representing $\operatorname{Lat^b}\Phi$.

(b) A direct consequence of (a) and Theorem III.5.2.

(c) From (a), [28] (or [29, Theorem 5.93]), and Theorem III.5.2. □

Example 4.5 (*Deformations of the braid arrangement*). These real affine hyperplane arrangements have recently been extensively studied. They are the affinographic representations $\mathcal{A}(\Phi)$ of certain integral additive gain graphs. To describe them let $\vec{K}_n$ denote the complete graph on node set $\{1, 2, \ldots, n\}$ with each edge oriented in the ascending direction. For $A \subseteq \mathbb{R}^+$, $A\vec{K}_n$ is the gain graph having edges of the form (e_{ij}, α) for $i < j$ and $\alpha \in A$. Arrangements $\mathcal{A}(A\vec{K}_n)$, where A is a finite subset of $\mathbb{Z}$, are called "deformations of the braid arrangement". Principal examples are the Shi arrangement, where $A = \{0, 1\}$, the Linial arrangement, where $A = \{1\}$; their "extensions", where $A = \{-l+1, = l+2, \ldots, l-1, l\}$ or $A = \{1, 2, \ldots, l\}$, respectively, where $l = 1, 2, 3, \ldots$; the composed partition arrangement, where $A = \{0, \pm 1\}$, and its extension the l-composed partition arrangement, where $A = \{0, \pm 1, \ldots, \pm l\}$. A rather weird case is our Example 4.6.

Athanasiadis [2] studied several arrangements of this type, epecially to compute their characteristic polynomials—which (as implied by Corollary 4.5(a) and Theorem II.5.3) are the balanced chromatic polynomials of the associated gain graphs. Athanasiadis' methods are generally based on gain graph coloring using a large cyclic gain group.

For a more thorough discussion of arrangements like these, with references, see [2] and [47, Examples 3.2 and 10.5–8].

4.2. **Characterization.**

How do we recognize lift matroids $L(\Omega)$ or their canonical representations? We repeat here a characterization of graphic lift matroids from the beginning of [46, Section 3].

Proposition 4.6. *Let M_0 be a finitary matroid on a point set E_0. M_0 contains a point e_0 such that M_0/e_0 is graphic if and only if $M_0 = L_0(\Omega)$ for some biased graph Ω with edge set $E_0 \setminus \{e_0\}$.*

Given M_0 such that M_0/e_0 is graphic, Ω can be any $(\Gamma, \mathcal{B})$ such that Γ is a graph with $G(\Gamma) \cong M_0/e_0$ and $\mathcal{B}$ consists of the circles whose closure does not contain e_0, or equivalently those that are dependent in M_0.

Proof. As mentioned in [46], this is implicit in [14, Section 6] as amplified in our Section II.3 near Theorem II.3.1. That the closure of a circle's containing e_0 is equivalent to the circle's being independent in M_0 follows from the representation of M_0 as $L_0(\Omega)$. □

For recognizing canonical representations there is another simple criterion.

Proposition 4.7. *A projective representation of $L(\Omega)$, $f : E \to \mathbb{P}$, is a canonical lift representation of Ω if and only if*

$$\bigcap_{C \notin \mathcal{B}} \overline{f(C)} \not\subseteq \bigcup_T \overline{f(T)},$$

where the range of C is all unbalanced circles, that of T is all maximal forests, and the overbar denotes projective closure.

Proof. If f is canonical, then $f(e_0) \in \bigcap_C \overline{f(C)} \setminus \bigcup_T \overline{f(T)}$.

If on the other hand $p_0 \in \bigcap_C \overline{f(C)} \setminus \bigcup_T \overline{f(T)}$, then extend f to E_0 via $f(e_0) = p_0$ and let M be the matroid on E_0 induced by f as extended. In $\mathbb{P}/p_0$, f induces a representation of M/e_0. Since $p_0 \in \overline{f(C)}$ if C is unbalanced, in that case $\operatorname{rk}_{M/e_0} C = \operatorname{rk}_M C - 1 = \#C - 1$. If C is balanced, take $e \in C$; then $f(e) \in \overline{f(C \setminus e)}$ because f represents $L(\Omega)$, but also $p_0 \notin \overline{f(C \setminus e)}$ because $C \setminus e$ extends to a maximal forest; hence, $p_0 \notin \overline{f(C)}$ and $\operatorname{rk}_{M/e_0} C = \operatorname{rk}_M C = \#C - 1$. As $p_0 \notin \overline{f(T)}$, $f'(T)$ is independent in $\mathbb{P}/p_0$. We conclude that $M/e_0 = G(\Gamma)$, whence M is a graphic lift $L(\Gamma, \mathcal{B}')$ for some bias. Moreover, $\mathcal{B}' \subseteq \mathcal{B}$ because $p_0 \in \overline{f(C)}$ for an unbalanced circle C, and $\mathcal{B} \subseteq \mathcal{B}'$ since $p_0 \notin \overline{f(C \setminus e)} = \overline{f(C)}$ for a balanced circle. M is finitary because it is projective. Therefore $M = L_0(\Omega)$; and the desired conclusion follows from Proposition 4.3. □

This proposition holds equally well for any representation $f : E \to E(M)$ of $L(\Omega)$ in any matroid M (as long as M is assumed finitary if Ω has infinite order).

4.3. **Switching and projective equivalence.**

Canonical lift representations of switching-equivalent F^+-gain graphs are projectively equivalent (see Formula (4.1)). As with bias representations (Section 2.3), the converse does not hold in general. Let Φ be an F^+-gain graph. If η is a switching function and $\alpha \in \operatorname{Aut} F$, then $z_{\Phi^{\eta\alpha}} \approx z_\Phi$ (that is, they are projectively equivalent). If $\langle \Phi \rangle$ is a contrabalanced circle or theta graph, we may have $\langle \Phi' \rangle = \langle \Phi \rangle$ and $z_{\Phi'} \approx z_\Phi$ but still $\Phi' \neq \Phi^{\eta\alpha}$. I suggest a conjecture similar to that for bias representations.

Conjecture 4.8. Let Φ and Φ' be unbalanced F^+-gain graphs of finite order with $\|\Phi'\| = \|\Phi\|$ and with $L(\Phi)$ connected. If $L(\Phi)$ has a $U_{2,4}$ minor, and if $x_{\Phi'} \approx x_\Phi$, then Φ' is obtained from Φ by switching and a field automorphism. If $L(\Phi)$ has no $U_{2,4}$ minor and $x_{\Phi'} \approx x_\Phi$, then Φ' need not be so obtained.

4.4. **Abstract gains and nonunique representation.**

As with Theorem 2.1, in Theorem 4.1 $\mathfrak{G}$ is a particular subgroup of F^+. If it is an abstract group, there may be many ways for it to embed in F^+. In order to understand projective uniqueness of canonical lift representations we have to know how uniqueness is affected by these different embeddings.

A *canonical lift representation* of a gain graph Φ with abstract gain group $\mathfrak{G}$ is a canonical lift representation of Φ over F, as in Theorem 4.1, obtained from any embedding $\varepsilon : \mathfrak{G} \hookrightarrow F^+$.

Proposition 4.9. *Suppose Φ is a gain graph whose gain group $\mathfrak{G}$ is generated by $\varphi(E_*)$. The canonical lift representations of Φ induced by different embeddings $\varepsilon_1, \varepsilon_2 : \mathfrak{G} \hookrightarrow F^+$ are projectively equivalent if and only if ε_1 and ε_2 are equivalent under an automorphism of F.*

Proof. Similar to that of Proposition 2.9. □

To take account of embeddings of the gain group into the field, we define Φ to have *projectively unique lift representation up to gain-group embedding* if every lift representation is projectively equivalent to a canonical lift representation with respect to some embedding of $\mathfrak{G}$ in F^+.

4.5. **Orthography and affinography.**

There is an affine version of the canonical lift representation of Φ. Take an affine space $\mathbb{A}$ over a skew field F, a hyperplane $\mathbb{A}_0$, and a vector $\xi \nparallel \mathbb{A}_0$. Represent $G(\Gamma)$ in $\mathbb{A}_0$ arbitrarily by $e \mapsto g(e)$ and add a certain multiple of ξ to get a vector $f(e)$ so that f is a representation of $L(\Phi)$. We call this an *orthographic representation* of $L(\Phi)$ because it begins with the represented graphic matroid $G(\Gamma)$ and adds a multiple of a transverse vector ξ (that we can think of as orthogonal). There are two difficulties in the construction of f: first, to represent $G(\Gamma)$ in $\mathbb{A}_0$, and second, to find the right multiple of ξ for each edge.

The first difficulty, though real since not all graphic matroids may have an affine representation over F, is understood. We know the conditions under which such a representation exists.

Proposition 4.10 (Crapo and Rota [12, Theorem 2, p. 16.10])**.** *Let Γ be a graph of finite order and F a skew field. $G(\Gamma)$ is representable in an affine space over F if and only if $\#F \geq \chi(\Gamma)$, the chromatic number.*

The second difficulty is our problem here. Assume that Γ is a link graph. (Loops and unbalanced edges do not have affine orthographic representations.) We want $f(e)$ to equal $g(e) + \varphi(e)\beta_e\xi$ where β_e depends only on the representation of $G(\Gamma)$ and the orientation used to compute $\varphi(e)$, but not on the gains themselves; thus we get a close analog of the lift representation. Our task is to determine β_e.

An orientation of a circle C is described in relation to a fixed orientation of Γ by the function τ_C defined for $e \in C$ by $\tau_C(e) = +1$ if e agrees with the direction of C and -1 if e opposes C. The two possible orientations of C give two choices for τ_C, which are negatives of each other.

Theorem 4.11. *Let Γ be a link graph. Suppose $g : E \to \mathbb{A}_0$ represents $G(\Gamma)$ in a hyperplane $\mathbb{A}_0$ of an affine space $\mathbb{A}$ over F, a skew field, and suppose ξ is a vector in $\mathbb{A}$ not parallel to $\mathbb{A}_0$. If $\operatorname{char} F \neq 2$, fix an orientation of Γ. Then there exist scalars β_e, $e \in E$, such that for every gain graph $\Phi = (\Gamma, \varphi, F^+)$ with gain group F^+, $f(e) = g(e) + \varphi(e)\beta_e\xi$ defines a representation f of $L(\Phi)$ in $\mathbb{A}$. (Here, except when $\operatorname{char} F = 2$, $\varphi(e)$ is to be calculated in the fixed orientation of e and β_e depends on the orientation.)*

The coefficients β_e can be calculated from g by choosing, for each circle C, affine dependence coefficients $\lambda_C(e)$ and an orientation τ_C of C so that for each edge $e \in E$, $\tau_C(e)/\lambda_C(e)$ is independent of the circle $C \ni e$, and setting $\beta_e = \tau_C(e)/\lambda_C(e)$.

By *affine dependence coefficients* we mean that the $\lambda_C(e)$ sum to 0, are not all zero, and satisfy $\sum_{e\in C} \lambda_C(e)g(e) = 0$.

Proof. First we set up machinery. The projective completion of $\mathbb{A}$, call it $\mathbb{P}$, is the projective quotient of a vector space Λ. (This means $\mathbb{P}$ is the set of lines of Λ, or equivalently of collinearity classes of nonzero vectors.) $\mathbb{A}_0$ corresponds to a hyperplane Λ_0 in Λ and the ideal hyperplane of $\mathbb{P}$, h_∞, corresponds to a hyperplane h_0. Take a linear form $\alpha : \Lambda \to F$

whose kernel is h_0 and let $h_1 = \{x \in \Lambda : \alpha(x) = 1\}$. Then $\mathbb{A}$ is naturally isomorphic to h_1, with $\mathbb{A}_0$ corresponding to $h_1 \cap \Lambda_0$. Thus the representation g can be regarded as a vector representation in Λ_0. The vector ξ in $\mathbb{A}$, not parallel to $\mathbb{A}_0$, is a vector in h_1, hence by translation in h_0, and $\xi \notin \Lambda_0$ because $\xi \nparallel \mathbb{A}_0$.

It is clear that we may assume $\Lambda = F^{1+N}$ with $\Lambda_0 = \{0\} \times F^N$. All vector representations of $G(\Gamma)$ are projectively equivalent. Consequently, we may assume the coordinate system chosen so that g represents $G(\Gamma)$ canonically in Λ_0: we mean that e is represented by $x(e) = \hat{w} - \hat{v} \in \Lambda_0$, where v and w are the endpoints of e and are labeled so that, in the fixed orientation, e is directed from v to w, and that $g(e)$ is the scalar multiple of $x(e)$ that lies in h_1. (In characteristic 2, any orientation of Γ will do—or none is needed—because $\hat{w} - \hat{v} = \hat{v} - \hat{w}$.)

Now we take the standard lift representation $z(e) = \hat{w} - \hat{v} + \varphi(e)\xi$, with $\varphi(e)$ calculated according to the fixed orientation, and project $z(e)$ into h_1. That is, we take $f(e) = \beta_e z(e)$ where β_e is the multiplier that carries $x(e)$ to $g(e)$. Thus $\beta_e = 1/\alpha(x(e))$. Then $f(e) = g(e) + \varphi(e)\beta_e\xi$ and, by Theorem 4.1, f is the desired representation of $L(\Phi)$ in $h_1 \cong \mathbb{A}$.

We omitted to verify that $\alpha(x(e)) \neq 0$. This is so because g is an affine representation: so $g(e) \in h_1$, and $g(e)$ is a multiple of $x(e)$.

That completes the proof of the first part of the theorem. What is missing is the ability to calculate β_e directly from the affine representation of $G(\Gamma)$, completely without Λ and h_1. The solution to that problem is based on the minimal affine dependencies of the points $g(e)$. For each circle $C = \{e_1, \ldots, e_l\}$ there are nonzero scalars λ_i such that $\sum_{i=1}^{l} \lambda_i g(e_i) = 0$ and $\sum_i \lambda_i = 0$. We need to choose β_e, independent of any gain function, so that the vectors $f(e) = g(e) + \varphi(e)\beta_e\xi$ are affinely dependent if and only if C is balanced. The only possible affine dependences of $f(e_1), \ldots, f(e_l)$ have coefficients that are a fixed multiple of the λ_i. Then $\sum_i \lambda_i f(e_i) = 0$ if and only if the coefficient of ξ is zero. That coefficient is

$$\sum_{i=1}^{l} \varphi(e_i)\lambda_i\beta_{e_i}.$$

Since $\varphi(C) = \sum_i \varphi(e_i)\tau_C(e_i)$, to make the coefficient of ξ zero for all balanced gains on C we must take $\lambda_i\beta_{e_i}$ to be a fixed nonzero multiple of $\tau_C(e_i)$. That is, $\beta_{e_i} = k_C\tau_C(e_i)/\lambda_i$. We define $\lambda_C(e_i) = \lambda_i/k_C$. The only problem is that β_{e_i} appears to depend on C. By the first part of the theorem, β_e is well defined, so it is possible to choose the scalars k_C so that, when $e \in C \cap C'$, then $\tau_C(e)/\lambda_C(e) = \tau_{C'}(e)/\lambda_{C'}(e)$. Thus we have proved the second half of the theorem. □

Because β_e depends only on g and the arbitrary orientation of Γ, not on Φ for parallel edges e and f we have $\beta_e = \beta_f$ if e and f are similarly oriented and $\beta_e = -\beta_f$ otherwise. Thus, if we choose to orient parallel edges in Γ similarly, β_e will depend only on the endpoints of e.

Corollary 4.12. *Under the hypotheses of Theorem 4.11, let $f(e_0)$ be the ideal point in $\mathbb{P}$, the projective completion of $\mathbb{A}$, that lies in the direction of ξ. This extension of f to E_0 is a projective representation of $L_0(\Phi)$.* □

This approach to orthographic representation is analytic. For a synthetic, coordinate-free treatment, see Part VI.

Example 4.6. Take $\Phi = \{-1, 1, 2\}\vec{K}_3$ with gain group F^+ (Figure 4.2(a)). This gain graph has edges γe_{ij} for $\gamma \in \{-1, 1, 2\}$ and $e_{ij} \in E(K_3)$ with $i < j$, the gain being $\varphi(\gamma e_{ij}) = \gamma$; that is, we assign gains in the upward orientation. We take this orientation for the arbitrary one of Γ in Theorem 4.11. We define $(-\gamma)e_{ji} = \gamma e_{ij}$ but oppositely oriented. The balanced circles,

$$C_1 = \{1e_{12}, 1e_{23}, 2e_{13}\}, \quad C_2 = \{(-1)e_{12}, 2e_{23}, 1e_{13}\}, \quad C_3 = \{2e_{12}, (-1)e_{23}, 1e_{13}\},$$

are the dependent triples in $L(\Phi)$ (assuming the characteristic is not too small), hence they generate the only collinear triples of points in the orthographic representation.

The affine space of the representation is $\mathbb{A} = \mathbb{A}^2(F)$; we take $\mathbb{A}_0 = \{(x_1, x_2) : x_2 = 0\}$ and $\xi = (0, 1)$. The representation of $G(K_3)$ in $\mathbb{A}_0$ consists of three distinct points with coordinates, let us say, $g(e_{ij}) = (\pi_{ij}, 0)$ for $i < j$. There is only one circle C in K_3, so to calculate the affine dependence coefficients $\lambda_C(e_{ij})$ we solve

$$\begin{bmatrix} 1 & 1 & 1 \\ \pi_{12} & \pi_{23} & \pi_{13} \end{bmatrix} \lambda_C^{\mathrm{T}} = 0,$$

whose solution is

$$\lambda_C = (\lambda_C(e_{12}), \lambda_C(e_{23}), \lambda_C(e_{13})) = (\pi_{13} - \pi_{23}, \pi_{12} - \pi_{13}, \pi_{23} - \pi_{12})$$

or any nonzero scalar multiple thereof. Orienting C in the direction $v_1v_2v_3v_1$, we have $\tau_C(e_{12}) = \tau_C(e_{23}) = -\tau_C(e_{13}) = 1$, so $\beta_{ij} = \tau_C(e_{ij})/\lambda_C(e_{ij})$ gives

$$\beta_{12} = (\pi_{13} - \pi_{23})^{-1}, \quad \beta_{23} = (\pi_{12} - \pi_{13})^{-1}, \quad \beta_{13} = (\pi_{12} - \pi_{23})^{-1}. \tag{4.2}$$

(We may write $\beta_{ij} = \beta(\gamma e_{ij})$, independent of γ, since parallel edges in Γ are similarly oriented.) The embedding of Φ is $f(\gamma e_{ij}) = (\pi_{ij}, \gamma\beta_{ij})$, so

$$\begin{aligned} f(\gamma e_{12}) &= (\pi_{12}, \gamma(\pi_{13} - \pi_{23})^{-1}), \\ f(\gamma e_{23}) &= (\pi_{23}, \gamma(\pi_{12} - \pi_{13})^{-1}), \\ f(\gamma e_{13}) &= (\pi_{13}, \gamma(\pi_{12} - \pi_{23})^{-1}). \end{aligned}$$

Now we choose a specific field, the finite field $\mathbb{F}_{53}$, so the gain group is $\mathbb{F}_{53}^+ = \mathbb{Z}_{53}$, and a particular embedding of $G(K_3)$ by taking $\pi_{12} = 0$, $\pi_{23} = 1$, and $\pi_{13} = 10$. Thus the orthographic points are:

	$\gamma = -1$	$\gamma = 1$	$\gamma = 2$
$f(\gamma e_{12})$	$(0, -6)$	$(0, 6)$	$(0, 12)$
$f(\gamma e_{23})$	$(1, 16)$	$(1, -16)$	$(1, 21)$
$f(\gamma e_{13})$	$(10, 1)$	$(10, -1)$	$(10, -2)$

Figure 4.2(b) shows the orthographic representation with its three 3-point lines corresponding to C_1, C_2, and C_3.

The orientation given to Γ in the theorem is arbitrary, but there is something to say about it when F is $\mathbb{R}$, or any ordered field.

Proposition 4.13. *If F is ordered, then there is a fixed orientation of Γ in Theorem 4.11 under which all β_e have the same sign. This orientation is acyclic and is unique up to reversal.*

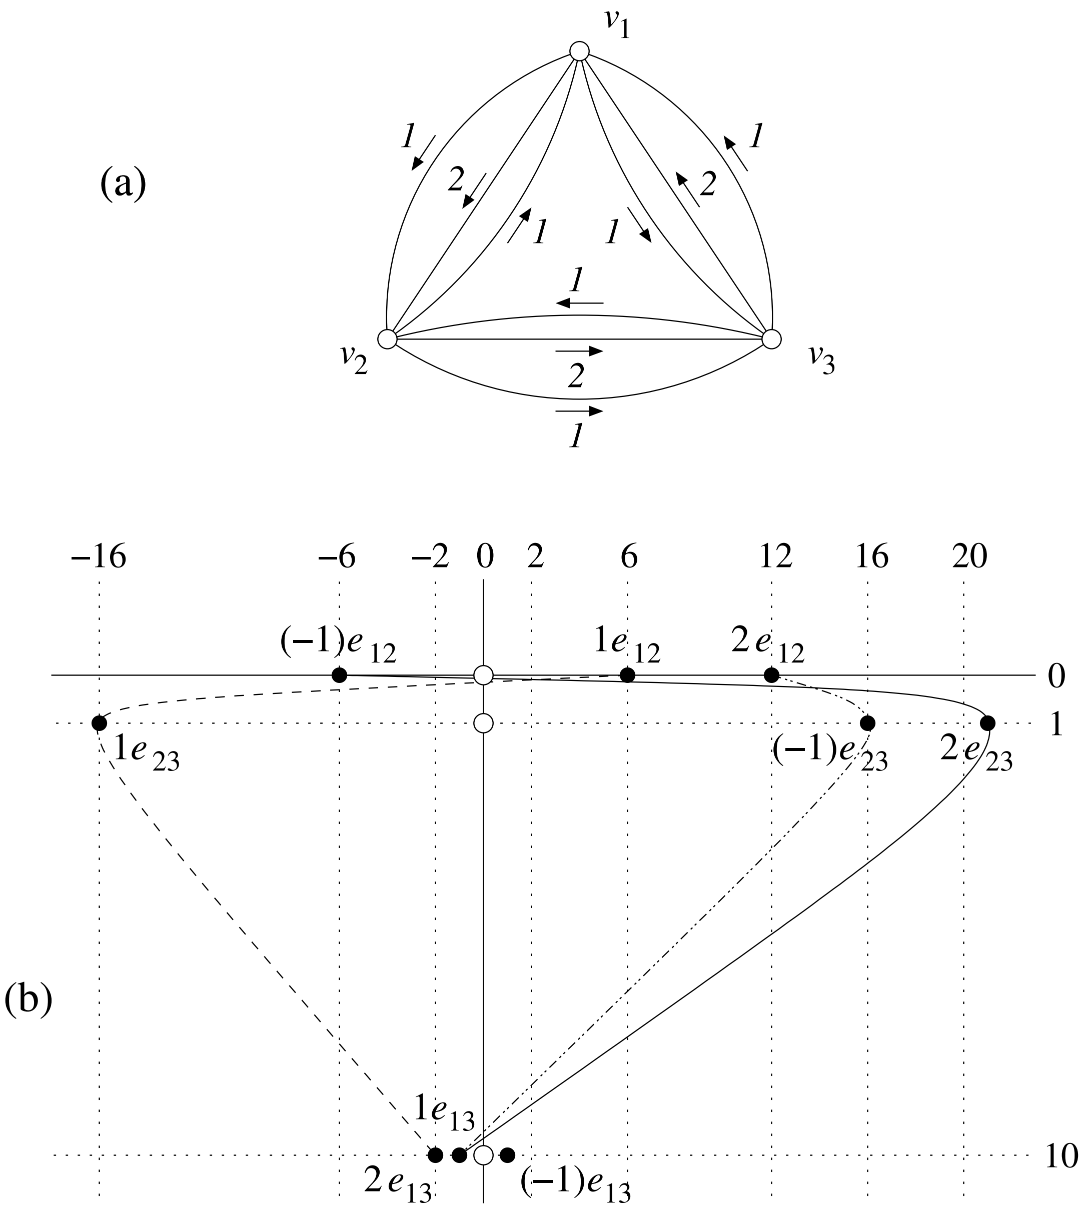

FIGURE 4.2. (a) The gain graph $\{-1, 1, 2\}\vec{K}_3$ of Examples 4.6 and 4.7. (b) The orthographic representation over $\mathbb{F}_{53}$ in Example 4.6. The hollow points represent $G(\|\Phi\|)$, or $G(K_3)$, in $\mathbb{A}_0$ (the solid vertical line). The solid points represent $L(\Phi)$ in $\mathbb{A}$ (the plane).

The proof is based on a series of reinterpretations, as in the proof of Theorem 4.11 but further extended. The hyperplane h_0 corresponds, in the dual space Λ_0^*, to a vector h_0^* in a region of the hyperplane arrangement $\mathcal{H}[\Gamma] = \{x_i = x_j : \text{there is an edge } ij \in E\}$. A region of $\mathcal{H}[\Gamma]$ corresponds to an acyclic orientation of Γ (see [17] or [18, Lemma 7.1]); negating the region corresponds to reversing the orientation. If we vary h_∞ continuously in $\mathbb{P}$ without passing over any point $g(e)$ (we call this *isotopy*), h_0^* varies in a region of $\mathcal{H}[\Gamma]$. Since h_∞

is not oriented, it corresponds to a pair of opposite vectors h_0^*; therefore an isotopy class of ideal hyperplanes corresponds to a converse pair of acyclic orientations of Γ. Furthermore, isotopy of h_∞ does not affect the combinatorial type of the affine representation of $G(\Gamma)$; in particular, the vertices of $\operatorname{conv} g(E)$ are unchanged.

If we choose one of the orientations corresponding to h_∞ as the fixed orientation in Theorem 4.11, we find that all $\alpha(x(e))$ have the same sign, whence all β_e have the same sign. The reason is that, when Γ is oriented in accordance with the acyclic orientation corresponding to the region that contains h_0^*, $\alpha(x(e)) > 0$ for all edges e. □

We can actually calculate the orientation in Proposition 4.13. If we calculate all β_e with respect to an arbitrary orientation and reverse the edges for which $\beta_e < 0$, we obtain the orientation with respect to which all $\beta_e > 0$. This method is explicit but complicated. It also seems to be using too much information about the representation. It is clear that the region in which h_0^* lies is determined in principle (up to negation) by the extreme points of $\operatorname{conv} g(E)$; but is this feasible in practice?

Problem 4.14. Is there a simple way to compute the acyclic orientations in Proposition 4.13 directly from the extreme points of $\operatorname{conv} g(E)$ in $\mathbb{A}$ or their oriented matroid?

Example 4.7. To illustrate Proposition 4.13 we take Example 4.6 in a different direction by choosing an ordered field, $F = \mathbb{R}$. The gain group is then $\mathbb{R}^+$. We use the same coordinates to represent $G(K_3)$: $\pi_{12} = 0$, $\pi_{23} = 1$, $\pi_{13} = 10$. What orientation of Γ makes all $\beta_e > 0$? (Since we oriented parallel edges similarly, we are really talking about an orientation of K_3.) We can deduce these from the signs of the β_{ij} in Equations 4.2: because we chose $\pi_{12} < \pi_{23} < \pi_{13}$, the signs are $+$, $-$, and $-$. Therefore, we must reverse the orientations of e_{23} and e_{13}, as shown in Figure 4.3(a). The new orientation is, as it ought to be, acyclic.

The new values of β_{ij} are

$$\beta_{12} = 1/9, \quad \beta_{32} = 1/10, \quad \beta_{31} = 1$$

and the orthographic points are

$$f(\gamma e_{12}) = \left(0, \tfrac{1}{9}\gamma\right), \quad f((-\gamma)e_{32}) = \left(1, \tfrac{1}{10}(-\gamma)\right), \quad f((-\gamma)e_{31}) = (10, (-\gamma)).$$

Figure 4.3(b) shows these points and the three 3-point lines associated with the balanced triangles C_1, C_2, and C_3.

The duals of orthographic representations are the affinographic hyperplane arrangements introduced at the end of Section 4.1. Under this duality, the dual of the orthographic point $f(e_0)$ is the ideal hyperplane h_∞, the lift vector ξ corresponds to the constant term of an affinographic hyperplane, and the orthographic base hyperplane $\mathbb{A}_0$ corresponds to the common line, $\operatorname{span}\{(1, 1, \ldots)\}$, of all homogeneous affinographic hyperplanes. In the orthographic representation all lines $f(E_0{:}\{u, v\})$ concur in $f(e_0)$ while in the projectivized affinographic arrangement $\mathcal{A}_{\mathbb{P}}(\Phi)$ all the colines $\bigcap \big\{h(e) : e \in E_0{:}\{u, v\}\big\}$ are contained in h_∞.

4.6. **Pythagorean representations.**

In conclusion we state a curious hyperplane representation for the complete lift matroid of Φ when Γ is a graph whose edges are links and the gain group is $\mathbb{R}^+$. We work in the Euclidean space $\mathbb{E}^d$ of any positive dimension. Let $N = \{v_1, \ldots, v_n\}$ and let $Q_1, \ldots, Q_n$ be

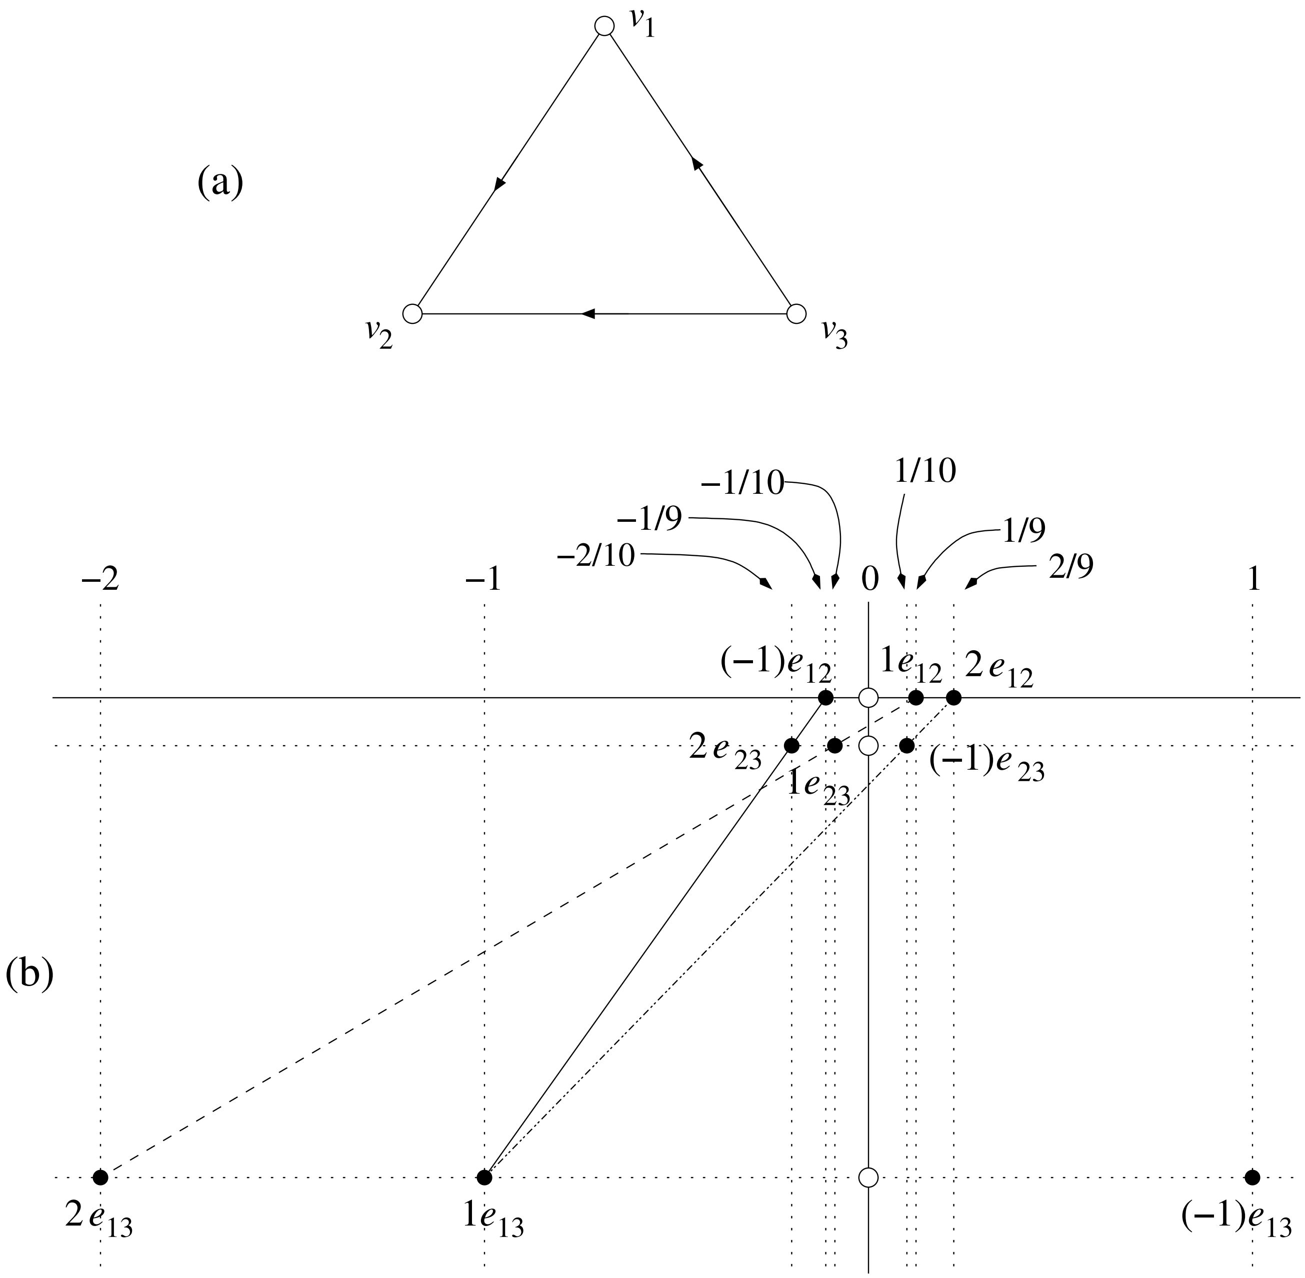

FIGURE 4.3. (a) The adjusted orientation of K_3 to make all $\beta_{ij} > 0$ in Example 4.7. (b) The real orthographic representation of $L(\Phi)$ from Example 4.7.

distinct points in $\mathbb{E}^d$. For $i \neq j$ and $P \in \mathbb{E}^d$, the *Pythagorean coordinate* of P from Q_i to Q_j is

$$\psi_{ij}(P) = \mathrm{dist}(P, Q_i)^2 - \mathrm{dist}(P, Q_j)^2.$$

For each edge e, with endpoints v_i and v_j, let

$$h(e) = \{P : \psi_{ij}(P) = \varphi(e; v_i, v_j)\}.$$

The set $\mathcal{H} = \{h(e) : e \in E\}$ is a family of Euclidean hyperplanes; $\mathcal{H}_{\mathbb{P}} = \mathcal{H} \cup \{h_\infty\}$ is a family of projective hyperplanes.

Theorem 4.15. *Let $n - c(\Phi) \geq d$. For $Q_1, \ldots, Q_n$ in general position, the set of flats of $\mathcal{H}_{\mathbb{P}}$, ordered by reverse inclusion, is isomorphic to the poset of elements of rank $\leq d$ in* $\operatorname{Lat} L_0(\Phi)$ *with a top element added (unless $n - c(\Phi) = d$ and Φ is balanced). The set of affine flats is isomorphic to the set of elements of rank $\leq d$ in* $\operatorname{Lat}^{\mathrm{b}} \Phi$.

A proof appears in [47] along with further structural information about $\mathcal{H}$ and $\mathcal{H}_{\mathbb{P}}$ and an exploration of the exact meaning of 'general position'.

The construction is metric: it can be carried out in any inner-product space F^d for gain graphs whose gain group is F^+.

This *Pythagorean representation* generalizes the canonical affine hyperplane lift representation. One way to assure general position is to choose $Q_1, \ldots, Q_n$ affinely independent. (See [47, Proposition 6.5] for a proof in the real case.) Let us take for Q_i the point in F^n whose coordinate vector equals $(1/\sqrt{2})b_i$, where b_i is the ith vector in an orthonormal basis B. A simple calculation shows that, since all $\|Q_i - Q_j\| = 1$, $\psi_{ij}(P) = c$ defines the hyperplane $h : x_j - x_i = c$. Thus for this choice of $Q_1, \ldots, Q_n$, the Pythagorean arrangement $\mathcal{H}$ is the canonical affine hyperplane lift representation of Φ.

5.1. **Whitney operations.**

Whitney's 2-isomorphism operations on a graph, which do not change the circles or, therefore, the polygon matroid, are:

(a) Identify two nodes in different components; and the inverse operation.
(b) Twist one side of a 2-separating node set.

Since these operations do not change circles or theta graphs regarded as edge sets, we can treat them as acting on biased graphs. Specifically, if $(\Gamma, \mathcal{B})$ is a biased graph and we apply Whitney operations to Γ, resulting in Γ', then $(\Gamma', \mathcal{B})$ is the resulting biased graph. It follows that Whitney operations do not change the lift or complete lift matroid. They can change the bias matroid; for instance, by a Whitney operation a contrabalanced handcuff can become disconnected, so no longer a bias circuit. However:

Theorem 5.1. *Let Ω' be a biased graph obtained from Ω by Whitney 2-isomorphism operations. Then Ω and Ω' have gains in the same groups; $G(\Omega)$ and $G(\Omega')$ are canonically representable over the same skew fields; and $L(\Omega) = L(\Omega')$ and $L_0(\Omega) = L_0(\Omega')$.*

Proof. The second conclusion follows from the first.

Consider the effect on gains φ for Ω of a single Whitney operation. Operations of type (a), retaining the same gains, obviously do not alter balance. Consider an operation (b) on Γ, twisting Γ_2 around the node pair u, v while leaving Γ_1, the remainder of the graph, untwisted, giving Γ'. We define φ', gains on Γ', by $\varphi'(e; x, y) = \varphi(e; x, y)$ if $e \in E_1$ and $\varphi(e; x, y)^{-1}$ if $e \in E_2$. Now examine a circle C with edges in both halves of Γ. Say C is the concatenation P_1P_2 in Γ, where P_i is the path of C in Γ_i. We know that $\varphi(P_1)\varphi(P_2) = 1$ and $\varphi'(P_2) = \varphi(P_2^{-1})$. Thus $\varphi'(P_1)\varphi'(P_2^{-1}) = 1$. But in Γ_2, $C = P_1P_2^{-1}$. Thus the gain of C remains the same, whence $\mathcal{B}(\Gamma', \Phi') = \mathcal{B}(\Gamma, \Phi)$. The theorem follows. □

The proof shows that we can think of Whitney operations as acting, not only on graphs and biased graphs, but also on gain graphs. Thus, we define two biased graphs, or two gain graphs, to be *2-isomorphic*, or *isomorphic up to Whitney operations*, if by applying Whitney operations they become isomorphic.

We can interpret the theorem as saying that, while $G(\Omega)$ is not determined by $G(\Gamma)$ and $\mathcal{B}$ alone, still the canonical representability of $G(\Omega)$ is so determined. If we know that $G(\Omega)$ has only canonical representations—as is the case with full biased graphs (Proposition 2.4) and thick biased graphs for which $G(\Omega) \neq L(\Omega)$ (see the discussion following Theorem 7.1)—then $G(\Gamma)$ and $\mathcal{B}$ determine all representations of $G(\Omega)$; Γ itself is not needed. This is odd, as gains of circles cannot even be defined directly on $G(\Gamma)$; one needs the graph.

5.2. **Separable biased graphs.**

The matroids of a separable biased graph, when not themselves separable according to Theorems II.2.8 and II.3.8, are 2-sums or parallel connections. Notation: $\Omega = \Omega_1 \cup_v \Omega_2$ means that $\Omega = \Omega_1 \cup \Omega_2$ and $\Omega_1 \cap \Omega_2 = \{v\}$. The 2-sum of matroids is written $M_1 \oplus_2 M_2$.

Theorem 5.2. *Suppose Ω_0 has a node v such that $\Omega_0 = \Omega_1 \cup_v \Omega_2$. Let h be an additional unbalanced edge at v and let $\Omega_i^h = \Omega_i \cup \{h\}$. Then $G(\Omega_0^h)$ is the parallel connection along h of $G(\Omega_1^h)$ and $G(\Omega_2^h)$, and $G(\Omega_0) = G(\Omega_1) \oplus_2 G(\Omega_2)$.*

Proof. The parallel connection along h of matroids M_1 and M_2 whose intersection is $\{h\}$ has for circuits the circuits of M_1, those of M_2, and the sets of the form $C_1 + C_2$ (this is set sum) where C_i is a circuit of M_i and $h \in C_1 \cap C_2$. In Ω_0^h there are three kinds of bias circuit: those of Ω_1^h, those of Ω_2^h, and those circuits C not contained in either E_1^h or E_2^h. Any such C has v as a cutpoint and is therefore a contrabalanced handcuff composed of unbalanced figures $D_1 \subseteq E_1$ and $D_2 \subseteq E_2$ and a connecting path P that contains v. Then $C = C_1 + C_2$ where $C_i = D_i \cup (P \cap E_i) \cup \{h\}$, a bias circuit of Ω_i.

Conversely, if C_i is a bias circuit of Ω_i^h and $h \in C_1 \cap C_2$, then C_2 and C_2 are contrabalanced handcuffs and so is $C_1 + C_2$. Therefore $G(\Omega_0^h)$ is the parallel connection.

The second part of the theorem follows because the 2-sum is the parallel connection with h deleted. □

Corollary 5.3. *$G(\Omega_0^h)$ is representable over F if and only if $G(\Omega_1^h)$ and $G(\Omega_2^h)$ are. Ω_0 has a canonical bias representation over F if and only if Ω_1 and Ω_2 do.*

Proof. The first is a general fact about parallel connection of matroids. The second is trivial. □

Theorem 5.4. *Suppose $\Omega = \Omega_1 \cup \Omega_2$ where $\Omega_1 \cap \Omega_2$ is void or a node. Then $L_0(\Omega)$ is the parallel connection along e_0 of $L_0(\Omega_1)$ and $L_0(\Omega_2)$, and $L(\Omega) = L(\Omega_1) \oplus_2 L(\Omega_2)$.*

Proof. Again we need only examine circuits of $L_0(\Omega)$ that do not lie in $(E_1)_0$ or $(E_2)_0$. By Theorem 5.1 we may assume that every block of Ω is a connected component. A circuit C not contained in $(E_1)_0$ or $(E_2)_0$ is therefore disconnected and consequently a union $D_1 \cup D_2$ where D_i is an unbalanced circle in E_i. Since each $D_i \cup \{e_0\}$ is a circuit in $L_0(\Omega_i)$, C has the form $C_1 + C_2$ where $e_0 \in C_1 \cap C_2$ and C_i is a circuit in $L_0(\Omega_i)$. Conversely, any $C_1 + C_2$ of that form is a lift circuit of Ω. □

Corollary 5.5. *$L_0(\Omega)$ is representable over F if and only if $L_0(\Omega_1)$ and $L_0(\Omega_2)$ are. Ω has a canonical lift representation over F if and only if Ω_1 and Ω_2 do.* □

6. Our matroids redefined by restricted general position

For the 'thick' representation theorem of Section 7 we need a new way to describe the bias and lift matroids. Suppose we have a biased graph Ω; let $H = \{h_v : v \in N\}$ consist of one half edge at each node, with H disjoint from E. Let $\Omega^\bullet = \Omega \cup H$ and $E^\bullet = E(\Omega^\bullet) = E \cup H$. (This is a variation on the usual meaning of $\Omega^\bullet$.) We write S_{vw} as shorthand for $S{:}\{v, w\}$ and S_v as shorthand for $S{:}\{v\}$. We consider finitary matroids M on point set $E^\bullet$ with various of the following properties: the general properties

- (a) $\operatorname{rk}_M S \leq \#N(S)$ for every finite $S \subseteq E^\bullet$ and $\operatorname{rk}_M\{\text{loose edges}\} = 0$;
- (b) $\operatorname{rk}_M S < \#N(S)$ for every finite balanced $S \subseteq E^\bullet$ that contains at least one ordinary edge;

the more special properties

- (a′) $\operatorname{rk}_M E^\bullet_{vw} \leq 2$, $\operatorname{rk}_M E^\bullet_v \leq 1$, and $\operatorname{rk}_M\{\text{loose edges}\} = 0$;
- (b′) Every balanced circle is dependent;
- (c) Every unbalanced circle or half edge in Ω is independent;

and the half-edge properties

- (g) H is independent;

(l) $M|H$ is a uniform matroid of rank 1;
(s) $M|H$ is simple.

Finally, we need the property of general position. $M' \geq M$, for matroids on the same points, means that every independent set of M is also independent in M'. Equivalently, $\mathrm{rk}_M S \leq \mathrm{rk}_{M'} S$ for every point set S. We say M' is *weaker* than M. Geometrically this means that M has more special position than does M'. The property we want is:

(w) M is the weakest matroid with whatever other properties are prescribed.

(It is not axiomatic that there is a unique such weakest matroid.)

The general properties are almost equivalent to the special ones.

Lemma 6.1. *(a′, g) $\implies$ (a′, s) $\implies$ (a).*

Proof. We prove (a′, s) $\implies$ (a). Since $\mathrm{rk}_M H_{vw} = 2 \geq \mathrm{rk}_M E^{\bullet}_{vw}$ by (a′), $E_{vw} \subseteq \mathrm{clos}_M H_{vw}$. Thus $E{:}X \subseteq \mathrm{clos}_M(H{:}X)$ for all $X \subseteq N$ of cardinality at least 2. For smaller X we apply (a′) directly. □

Lemma 6.2. *(b′) $\iff$ (b).*

Proof. We prove (b′) $\implies$ (b). Let S be a connected, balanced edge set and T a basis for $M|S$. If $\mathrm{rk}_M S \geq \#N(S) > 0$, then T contains a circle C. C is independent because it lies in T but is dependent because S is balanced. This is a contradiction.

If S is balanced with components $S_1, \ldots, S_k$, then $\mathrm{rk}_M S \leq \mathrm{rk}_M S_1 + \cdots + \mathrm{rk}_M S_k < \#N(S_1) + \cdots + \#N(S_k) = \#N(S)$. □

Our main results are characterizations of G and L, mostly in terms of (w) but in one case with (g) instead of (w).

Proposition 6.3. *(a, b, l, w) $\iff$ $M = L(\Omega^{\bullet})$. That is, there is a unique weakest matroid satisfying (a, b, l), and it is $L(\Omega^{\bullet})$.*

Proof. Let M satisfy (a, b, l). Since M and $L(\Omega^{\bullet})$ are finitary it suffices to treat finite graphs. We write $S_1, \ldots, S_k$ for the components of an edge set $S \subseteq E^{\bullet}$, the first j being the unbalanced ones (so $0 \leq j \leq k$), and $n_i = \#N(S_i)$. If S is balanced ($j = 0$), then

$$\mathrm{rk}_L S = \sum_i \mathrm{rk}_L S_i = \sum_i (n_i - 1)$$

and from (b),

$$\mathrm{rk}_M S \leq \sum_i \mathrm{rk}_M S_i \leq \sum_i (n_i - 1),$$

so that $\mathrm{rk}_L S \geq \mathrm{rk}_M S$. If S is unbalanced, then

$$\mathrm{rk}_L S = \sum_i (n_i - 1) + 1 = \sum_i n_i - k + 1.$$

Let us consider an unbalanced component S_i. Either $\mathrm{rk}_M S_i < n_i$, or S_i spans $E^{\bullet}{:}N(S_i)$ (by (a)) so has a basis B_i that includes a half edge h_{v_i} where $v_i \in N(S_i)$. Suppose the latter applies to components $S_1, \dots, S_l$ with $l > 0$. Then setting $S'_i = S_i \cup \{h_{v_i}\}$,

$$\begin{aligned}
\mathrm{rk}_M(S'_1 \cup \cdots \cup S'_l) &= \mathrm{rk}_M(B_1 \cup \cdots \cup B_l) \\
&\le \sum_{i=1}^{l} \mathrm{rk}_M(B_i \setminus H) + \mathrm{rk}_M\{h_{v_1}, \dots, h_{v_l}\} \\
&= \sum_{i=1}^{l} (n_i - 1) + 1.
\end{aligned}$$

Consequently,

$$\begin{aligned}
\mathrm{rk}_M S &\le \sum_{i=1}^{l} (n_i - 1) + 1 + \sum_{i=l+1}^{j} (n_i - 1) + \sum_{i=j+1}^{k} (n_i - 1) \\
&= \mathrm{rk}_L S.
\end{aligned}$$

This formula applies as well when $l = 0$. We have shown that $\mathrm{rk}_M \le \mathrm{rk}_L$; this establishes that L is the unique weakest matroid that satisfies (a, b, l). □

This proposition is really about $L_0(\Omega)$. Since H is contained in an atom, identifying it to a single point e_0 converts $L(\Omega^\bullet)$ to $L_0(\Omega)$.

Proposition 6.4. *(a′, b′, s, w) $\iff$ (a, b, w) $\iff$ $M = G(\Omega^\bullet)$. That is, there is a unique weakest matroid satisfying (a′, b′, s) or (a, b), and it is $G(\Omega^\bullet)$.*

Proof. Again, it suffices to treat finite graphs. As in the preceding proof, but more simply,

$$\mathrm{rk}_M S \le \sum_{i=1}^{k} \mathrm{rk}_M S_i \le n_1 + \cdots + n_j + (n_{j+1} - 1) + \cdots + (n_k - 1)$$

by (a, b), and this $= \mathrm{rk}_G S$. □

A different kind of characterization replaces (w) by an explicit prescription of $M|H$.

Proposition 6.5. *(a′, b′, c, g) $\iff$ (a, b, c, g) $\iff$ $M = G(\Omega^\bullet)$. That is, $G(\Omega^\bullet)$ is the unique matroid satisfying (a, b, c) in which H is independent.*

Proof. (a′, g) imply that M is a frame matroid with special basis H. (See [44] for frame matroids. The defining property is simply that the lines generated by H contain all points.) By [44, Theorem 1], $M = G(\Omega')$ where Ω' is a biased graph on the edge set $E^\bullet$. Since (b′) requires all Ω-balanced circles to be dependent in M, they are balanced in Ω'. That is, $\mathcal{B}' \supseteq \mathcal{B}$. Contrariwise, by (c) a circle $C \notin \mathcal{B}$ is independent, so $C \notin \mathcal{B}$'. Therefore $\mathcal{B}' = \mathcal{B}$, and the result is proved. □

7. Thick biased graphs and intermediate-matroid representation

It is a difficult problem to characterize all bias and lift representations of a biased graph, but multiple edges make it easier and for a graph with enough multiple edges we can give a complete solution. We call a biased graph Ω *thick* if, whenever v and w are adjacent (written $v \sim w$), then $E{:}\{v, w\}$ is unbalanced. (Equivalently, $\mathrm{rk}(E{:}\{v, w\}) = 2$ in $G(\Omega)$ and $L(\Omega)$.) For the matroids of thick biased graphs, there are no vector representations other than the canonical lift and bias representations of gain graphs with the given bias.

Our theorem treats more than bias and lift representations. An *intermediate matroid* on Ω is a matroid M with point set $E(\Omega)$ such that $G(\Omega) \geq M \geq L(\Omega)$ (as defined at the start of Section 6).

Theorem 7.1 (Thick representation). *Let Ω be a biased graph of order at least 3 that contains as a spanning subgraph a thick, 2-connected biased graph. Suppose M is a finitary intermediate matroid on Ω which is representable in some projective space. Then $M = L(\Omega)$ or $G(\Omega)$ and every representation of M, in any projective space, is a canonical lift or bias representation. (If $L(\Omega) = G(\Omega)$, M may have both kinds of representation.)*

We might call this a 'semi-unique representation' theorem. We shall analyze its significance after giving the proof.

Proof. We may assume that Ω is simply biased; that is, it has no loose edges, balanced loops or digons, or pairs of unbalanced edges with the same supporting node. We also may replace any unbalanced loops by half edges. We write $E_{vw} = E{:}\{v, w\}$.

First we treat the case in which Ω itself is thick.

We suppose M represented in a projective space $\mathbb{P}$ by a mapping $\theta : E \to \mathbb{P}$. Let $\hat{e} = \theta(e)$, the projective point representing edge e; and for $S \subseteq E$, let $\hat{S} = \{\hat{e} : e \in S\}$. When $v \sim w$, let L_{vw} be the line determined by $\hat{E}_{vw}$. (Otherwise L_{vw} is undefined.) We write $\overline{X}$ for the projective span of $X \subseteq \mathbb{P}$; thus $L_{vw} = \overline{\hat{E}_{vw}}$. A frequently used fact is that, if $\operatorname{rk}_G S = \operatorname{rk}_L S$, then $\operatorname{rk}_M S =$ their common value.

First we note that $L_{vw} \cap \hat{E}_* \subseteq \hat{E}_{vw}$, because E_{vw} has rank 2 in G and L, hence also in M, but adding any link raises the rank. For the same reason no two lines L_{vw} coincide.

Next, we show that node points are well defined. (It is here that we require $n \geq 3$.)

Lemma 7.2. *For each $v \in N$, there is a unique point $p_v \in \mathbb{P}$ such that for any $w \sim v$, $p_v \in L_{vw}$. If v supports an unbalanced edge h_v, then $p_v = \hat{h}_v$; otherwise, $p_v \notin \hat{E}$. Furthermore, for any two neighbors w, x of v, $L_{vw} \cap L_{vx} = \{p_v\}$.*

Proof. Since $\operatorname{rk}_M(E_{vw} \cup E_{vx}) = 3$, L_{vw} and L_{vx} are coplanar; therefore $L_{vw} \cap L_{vx}$ is a point p_{wx}. If h_v exists, $p_{wx} = \hat{h}_v$.

Now take $y \sim v$. Since $\operatorname{rk}_M(E_{vw} \cup E_{vx} \cup E_{vy}) = 4$, L_{vw}, L_{vx}, and L_{vy} are noncoplanar. The only way p_{wx}, p_{wy}, and p_{xy} can all exist is if they are equal. It follows that all p_{wx}, for $w, x \sim v$, are the same point, which we call p_v. □

Lemma 7.3. *Let C be a circle in Ω of length at least 3. Either all p_v for $v \in N(C)$ are equal, or all are projectively independent.*

Proof. Let $N(C) = (v_1 v_2 \cdots v_l)$ in cyclic order around C. Let $L_i = L_{i-1,i}$ (with subscripts modulo l). Let us call L_i *loose* if $p_{i-1} = p_i$ and *stiff* if $p_{i-1} \neq p_i$. If L_i is loose, let $q_i \in L_i \setminus p_i$. If L_i is stiff, call i a *core index*; let I be the set of core indices. If there are no core indices, it is because all p_i are equal.

Suppose a core index exists. Let $S = E_{01} \cup E_{12} \cup \cdots \cup E_{l-1,l}$. Since $L_i = \overline{p_{i-1} p_i}$ if stiff and $\overline{p_i q_i}$ if loose, and since $\{p_i : 1 \leq i \leq l\} = \{p_i : i \in I\}$, we have

$$\hat{S} \subseteq \overline{\{p_i : i \in I\} \cup \{q_j : j \notin I\}}.$$

The rank of $\hat{S}$ is $\operatorname{rk}_M S = l$, since S has rank l in G and L. The rank of the right-hand subspace is at most l because it is generated by l points. Thus if all L_i are stiff, $p_1, \ldots, p_l$

are independent. If there is a loose line, say L_1, let $T = E_{12} \cup \cdots \cup E_{l-1,l}$. Then $\mathrm{rk}_M T = l$. But

$$\hat{T} \subseteq L_2 \cup L_3 \cup \cdots \cup L_l \subseteq \overline{\{p_i : i \in I\} \cup \{q_j : j \notin I \text{ and } j \neq 1\}}.$$

The right-hand generating set has only $l-1$ points, so its rank is less than l. This is a contradiction. Hence all lines are stiff and all node points p_i are projectively independent. $\square$

Now let us consider the possibility that one circle, C_0, has all node points equal and another, C, has its node points independent. By Menger's theorem we can embed $C \cup C_0$ in a 2-connected finite subgraph of Ω. By Tutte's path theorem [33, Theorem 4.34] in that subgraph, there is a chain of circles, $C_0, C_1, \dots, C_m = C$, such that C_{i-1} and C_i share an edge. This implies that C_{i-1} and C_i both have all node points equal or both do not. The contradiction is obvious.

Therefore two cases are possible: all node points p_v for $v \in N$ are equal, or they are not. In the latter case, suppose there were a finite set X of nodes whose corresponding node points had rank less than $\#X$. Then X can be extended to a finite set Y such that $\Omega_1 = \Omega{:}Y$ is 2-connected and $\{p_v : v \in Y\}$ has rank less than $\#Y$. However, in Ω_1,

$$\hat{E}_1 \subseteq \bigcup_{v \sim w} L_{vw} \subseteq \overline{\{p_v : v \in Y\}}.$$

The ranks are, on the left, $\mathrm{rk}_M(E_1) = \#Y$ because $L(\Omega_1)$ and $G(\Omega_1)$ have rank $\#Y$, and on the right, at most $\#Y$ because of the number of p_v's. Consequently, the points p_v for $v \in Y$ are projectively independent. It follows that $\{p_v : v \in N\}$ is independent.

That concludes the first part of the proof. We now must show that if all $p_v = p_0$, then $M = L(\Omega)$ and the representation is a canonical lift representation, while if the p_v are independent, then $M = G(\Omega)$ and the representation is a canonical bias representation.

Case 1: All $p_v = p_0$. That is, there is a point $p_0 \in \mathbb{P}$ that lies on all lines L_{vw}. We must prove that $M = L$. In fact, we prove more. Let M_0 be the one-point extension of M on the set $E_0 = E \cup \{e_0\}$, such that M_0 is represented by $\hat{E} \cup \{p_0\}$, extending the representation of M and with p_0 representing e_0. We show that $M_0 = L_0(\Omega)$.

First we demonstrate that $M_0 \geq L_0$. An independent set in $L_0 \setminus e_0 = L$ is necessarily independent in M. A finite independent set in L_0 that contains e_0 is a forest F together with e_0. Let $F_0 = F \cup \{e_0\} = \{e_0, e_1, \dots, e_k\}$. Since $\hat{e}_0 = p_0$, $\overline{\hat{F}_0}$ contains all the lines L_{vw} for nodes v and w that are adjacent in F. Thus $\mathrm{clos}_{M_0}(F_0)$ contains all edges parallel to those of F. The rank of this set of edges is $k+1$ in L and therefore is no less in M. Thus $\mathrm{rk}_{M_0} F_0 \geq k+1 = \#F_0$, and it follows that F_0 is independent in M_0. We conclude that $M_0 \geq L_0$; consequently, also $M_0/e_0 \geq L_0/e_0 = G(\|\Omega\|)$.

Now we show that $M_0/e_0 = L_0/e_0$. We must prove that a circle C in $\|\Omega\|$ is dependent in M_0/e_0. If C (of length l) is balanced, then it is dependent in G, hence in M, and consequently in M/e_0. If C is unbalanced, it is independent in L so $\mathrm{rk}_M(C) = l$. Take e'_{vw} parallel to an edge $e_{vw} \in C$, so $L_{vw} = \overline{\hat{e}_{vw}\hat{e}'_{vw}}$. We know $\mathrm{rk}_M(C) = \mathrm{rk}_M(C \cup e'_{vw}) = l$, therefore $\overline{\hat{C}} = \overline{\hat{C} \cup \hat{e}'_{vw}} \supseteq L_{vw} \ni p_0$. Thus $\mathrm{rk}(\hat{C} \cup p_0) = l$, which means $\mathrm{rk}_{M_0}(C \cup e_0) = l$; thus C has rank $l-1$ in M_0/e_0. That is, C is dependent. If e_{vw} has no parallel in E, then a half edge h_v (or h_w) is in E; then $C \cup h_v$ is a circuit. Since h_v is parallel to e_0 in L_0, again C is dependent in M_0/e_0.

We have just seen not only that $M_0/e_0 = G(\|\Omega\|)$ but also that a circle is dependent in M_0 if and only if it is balanced. It follows from Proposition 4.6 that $M_0 = L_0(\Omega)$.

Case 2: $\{p_v : v \in N\}$ is an independent set. By Lemma 7.4, $M^{\bullet} = G(\Omega^{\bullet})$.

We conclude that the theorem holds when Ω is thick. For if all $p_v = p_0$, then $M = L(\Omega)$ and the original representation extends to a representation of $L_0(\Omega)$, hence is a canonical lift representation. If all p_v are independent, then $M = G(\Omega)$ and the representation extends to one of $G(\Omega^{\bullet})$, hence is a canonical bias representation.

The lemma we apply to Case 2 is very general. Suppose a matroid M on $E(\Omega)$ is represented in a projective space $\mathbb{P}$ so that there exist points $p_v \in \mathbb{P}$, corresponding to the nodes, that are independent and such that, if $e \in E_{vw}$, then $\hat{e} \in \overline{p_v p_w}$. Then we say the representation has *a complete set of node points.*

Lemma 7.4. *Let Ω be a biased graph with $n \geq 3$ and M an intermediate matroid on Ω that is represented in some projective space with a complete set of node points. Then the representation is a canonical bias representation of Ω.*

Proof. Let $M^{\bullet}$ be M together with the node points p_v. Then $M^{\bullet}$ satisfies (a, g) of Section 6 because of the complete set of node points. It satisfies (b, c) because it is an intermediate matroid. By Proposition 6.5 it is $G(\Omega^{\bullet})$. It follows that $M = G(\Omega)$. □

The proof of the complete Theorem 7.1 depends on the generality of Lemma 7.4. Let Ω be as in the theorem and let Ω_t be a maximal thick subgraph of Ω. Thus Ω_t is 2-connected and spanning. By the thick case applied to Ω_t, either all edge lines are concurrent or there is a complete set of node points.

In the latter case, let the node points of Ω_t be p_v for $v \in N$ and write $\hat{X} = \{p_v : v \in X\}$. We know the node points of Ω_t are independent. We have to prove they are node points of Ω: that is, if e_{vw} is a link without a parallel edge, then $\hat{e}_{vw} \in \overline{p_v p_w}$. There are internally disjoint paths P_1 and P_2 in Ω_t from v to w. Let e_1 be parallel to an edge in P_1; then $P_1 \cup \{e_1\}$ is unbalanced. Either $P \cup \{e_1, e_{vw}\}$ contains a balanced circle on e_{vw}, which is a circuit in L and G, hence in M, or $P \cup \{e_1, e_{vw}\}$ is a contrabalanced theta graph, which is a circuit in L and G, thus in M. In either case, $e_{vw} \in \operatorname{clos}_M(P \cup \{e_1\})$. Therefore, $\hat{e}_{vw} \in \overline{\hat{N}(P_1)}$. Similarly, $\hat{e}_{vw} \in \overline{\hat{N}(P_2)}$. We deduce that $\hat{e}_{vw} \in \overline{\hat{N}(P_1)} \cap \overline{\hat{N}(P_2)} = \overline{p_v p_w}$. By Lemma 7.4, $M = G(\Omega)$ and the representation is a canonical bias representation.

What remains is the case in which all edge lines are concurrent at a point p_0. Note that an edge line L_{vw} exists if and only if $v \sim w$ and E_{vw} is unbalanced; and by Lemma 7.2 applied to Ω_t, if e is a half edge at v, then $\hat{e} = p_0$.

In order to prove that θ is a canonical lift representation we must show that it extends to a representation of $L_0(\Omega)$. We do this by extending M to a matroid $M^{\bullet}$ in the sense of Section 6, that is, $E^{\bullet} = E \cup H$ where $H = \{h_v : v \in N\}$ consists of one half edge for each node and is disjoint from E. Extending θ to $E^{\bullet}$ by $\theta(h_v) = p_0$ defines $M^{\bullet}$ as the matroid represented by the extended θ; our task is to prove that $M^{\bullet} = L(\Omega^{\bullet})$.

We do so by means of Proposition 6.3. Properties (b′) and (l) are obvious. That $M^{\bullet}$ is finitary follows from the finitarity of any projective dependence matroid. It remains to establish (a).

For finite $S \subseteq E$, (a) is valid because $\operatorname{rk}_M S \leq \operatorname{rk}_G S$. We therefore consider a set $S \subseteq E^{\bullet}$ such that $S \cap H \neq \varnothing$. The first step depends on the nature of $S \cap E$.

If $S \cap E$ is balanced, then $\operatorname{rk}_{M^{\bullet}} S \leq \operatorname{rk}_M(S \cap E) + \operatorname{rk}_{M^{\bullet}}(S \cap H) = \operatorname{rk}_M(S \cap E) + 1 < \#N(S \cap E) + 1$; so (a) is satisfied.

Suppose $R = S \cap E$ is unbalanced. If $p_0 \in \overline{\hat{R}}$, then $\mathrm{clos}_{M^\bullet} R \supseteq S$ so $\mathrm{rk}_{M^\bullet} S = \mathrm{rk}_M R$ and (a) is satisfied. Thus we must show that $p_0 \in \overline{\hat{C}}$ when C is an unbalanced circle. There are three cases.

Case 1: If $E{:}N(C)$ contains a double link $\{e, e'\}$, then

$$p_0 \in \overline{\hat{C} \cup \{\hat{e}, \hat{e}'\}},$$

so $\mathrm{rk}_{M^\bullet} C^\bullet \leq \mathrm{rk}_M(C \cup \{e, e'\}) \leq \mathrm{rk}_G(C \cup \{e, e'\}) \leq \#N(C) = \mathrm{rk}_M C$. Therefore $\hat{C}$ spans $\hat{C} \cup \{\hat{e}, \hat{e}'\}$; consequently $p_0 \in \overline{\hat{C}}$.

Case 2: If there is a half edge e' of Ω such that $C \cup e'$ is a circuit in M, then $p_0 \in \overline{\hat{C}}$ because $p_0 = \theta(e')$ by Lemma 7.2. This applies in particular if $N(C)$ supports a half edge in Ω.

Case 3: If neither Case 1 nor Case 2 applies, choose two nodes $v, w \in N(C)$. By Menger's theorem there are internally disjoint paths $vv' \cdots w$ and $ww' \cdots v$ in Ω_t. We focus on $E_{vv'}$. It contains a pair of parallel links, $\{e, e'\}$, or a link e and a half edge e' which (since we are not in Case 2) is at v'. Again since we are not in Case 2, in both of these subcases $C \cup \{e, e'\}$ is a circuit in M. But also, $p_0 \in \overline{\hat{e}\hat{e}'}$. Therefore, $p_0 \in \overline{\hat{C} \cup \hat{e}}$. Similarly, $p_0 \in \overline{\hat{C} \cup \hat{f}}$ if $f{:}ww'$ is a link. Since $\mathrm{rk}_M(C \cup e) = \mathrm{rk}_M(C \cup f) = \#N(C) + 1$ and $\mathrm{rk}_M(C \cup \{e, f\}) = \#N(C \cup \{e, f\}) = \#N(C) + 2$ (all because $C \cup e$, $C \cup f$, and $C \cup \{e, f\}$ are connected), by the submodular law

$$\overline{\hat{C} \cup \hat{e}} \cap \overline{\hat{C} \cup \hat{f}} = \overline{\hat{C}}.$$

Because p_0 belongs to both terms on the left side, $p_0 \in \overline{\hat{C}}$.

Now we demonstrate that $M^\bullet \geq L(\Omega^\bullet)$. Since $M^\bullet$ is finitary, it suffices to show that a finite independent set S in $L(\Omega^\bullet)$ is independent in $M^\bullet$. If $S \subseteq E$, S is independent in M by the intermediacy of M. Otherwise, S consists of a half edge $h \in H$ and a finite forest $T = S \setminus h$ in Ω.

Suppose $p_0 \in \overline{\hat{T}}$. If T does not connect N, there is a link $e \in E(\Omega_t)$ whose endpoints v and w are not connected by T. E_{vw} contains another edge e', and since $p_0\hat{e}\hat{e}'$ or $p_0\hat{e}'$ is a circuit and $p_0 \in \overline{\hat{T}}$, $\hat{e}' \in \overline{\hat{T} \cup \hat{e}}$. Thus, $e' \in \mathrm{clos}_M(T \cup e)$. However, $T \cup e$ is balanced and $T \cup \{e, e'\}$ is not, so they cannot have the same rank in M. This contradiction shows that T must be a spanning tree of Ω. In fact, we may assume that n is finite and no forest U that is not a spanning tree has p_0 contained in $\overline{\hat{U}}$. Hence, $\hat{T} \cup p_0$ is a circuit in $M^\bullet$.

Furthermore, if there is an edge $e \in T \cap E(\Omega_t)$, then $p_0 \in \overline{\hat{T}}$ implies $E{:}N(e) \subseteq \mathrm{clos}_{M^\bullet} T$, so $\mathrm{clos}_M T$ is unbalanced. This is again a contradiction. Therefore, $T \cap E(\Omega_t) = \varnothing$.

Since $n \geq 3$, T has an end node z, which is adjacent to a node y by an edge $f \in T$. Because Ω_t connects y to z but $f \notin E(\Omega_t)$, y is adjacent to some $x \neq z$ in Ω_t. Let $e, e' \in E_{xy}$. Now,

$$\hat{f} \in \overline{\hat{T} \setminus \hat{f} \cup p_0} \subseteq \overline{\hat{T} \setminus \hat{f} \cup \{\hat{e}, \hat{e}'\}},$$

so $f \in \mathrm{clos}_M(T \setminus f \cup \{e, e'\})$. However, this is impossible, as $T \setminus f \cup \{e, e'\} \subseteq E{:}(N \setminus z)$ and $\mathrm{clos}_M(E{:}(N \setminus z)) \subseteq E{:}(N \setminus z)$ because M is intermediate. We conclude that $p_0 \notin \overline{\hat{T}}$ and consequently $M^\bullet \geq L(\Omega^\bullet)$.

We have shown that $M^\bullet$ satisfies (a, b′, l) of Section 6 with respect to $\Omega^\bullet$. The weakest such matroid is $L(\Omega^\bullet)$. On the other hand, $M^\bullet \geq L(\Omega^\bullet)$. It follows that $M^\bullet = L(\Omega^\bullet)$, whence $M = L(\Omega)$ and θ is a canonical lift representation of Ω. $\square$

Let us discuss the meaning of Theorem 7.1.

First of all, it is a partial unique-representation theorem for $G(\Omega)$ and $L(\Omega)$ and it proves Conjectures II.2.15 and II.3.14 for sufficiently thick graphs. Suppose the bias and lift matroids are not equal, as for instance when Ω has disjoint unbalanced circles. Then for each gain group $\mathfrak{G}$ of Ω, $G(\Omega)$ has a unique representation (up to gain-group embedding) over each skew field F for which $\mathfrak{G} \hookrightarrow F^*$, and $L(\Omega)$ has a unique representation (up to gain-group embedding) over each skew field F for which $\mathfrak{G} \hookrightarrow F^+$. If $\mathfrak{G}$ is a subgroup of both F^* and F^+ (as is $\mathbb{Z}$, for instance, when char $F = 0$; but this can only happen when F is infinite), then $G(\Omega)$ and $L(\Omega)$ both have representations over F, canonically induced by the $\mathfrak{G}$-gains. These representations are, of course, inequivalent since $G(\Omega) \neq L(\Omega)$. Moreover, there are no other representations. (I am ignoring here the possibility that Ω has more than one essentially different $\mathfrak{G}$-gain function. That will produce additional representations.)

Perhaps more remarkable is that, when $G(\Omega)$ and $L(\Omega)$ are equal, this one matroid has two different kinds of representation. Again, each gain group $\mathfrak{G} \leq F^*$ or F^+ produces a canonical bias or lift representation over F for each embedding $\mathfrak{G} \hookrightarrow F^*$ or $\mathfrak{G} \hookrightarrow F^+$. If $\mathfrak{G} \leq F^*$ and F^+, the one matroid has (at least) two representations canonically induced by the $\mathfrak{G}$-gains. These representations are inequivalent (since $n \geq 3$): we see this from the fact that the canonical bias representation extends to node points p_i that are all distinct, so the edge lines L_{ij} are not concurrent, but in the canonical lift representation the edge lines are concurrent at p_0.

This double representability when $G(\Omega) = L(\Omega)$ is even more striking when Ω has unique gains. (I mean that any two gain functions for Ω are switching equivalent after appropriately switching and cutting down the gain group. Biased graphs with this property include those of the group expansions of nontrivial 2-connected graphs; see Theorem V.2.1.) Then $G(\Omega)$ has exactly two inequivalent representations over any skew field that contains the gain group uniquely both additively and multiplicatively.

We have already seen in Examples 2.1, 2.2, 4.1, and 4.2 applications of Theorem 7.1 to limit the variety of representations of a bias or lift matroid. In the former two all representations are canonical bias, in the latter all are canonical lift. Modified graphs exemplify the existence of both types of representation of the same matroid.

Example 7.1 (*Example 2.1 continued*). Let Φ_1 be Φ of Example 2.1 without h_1 and $(-1)e_{22}$. Most of Example 2.1 applies but there is a big difference: $G(\Phi_1)$ has noncanonical bias representations. These arise from Theorem 7.1 and the fact that $G(\Phi_1) = L(\Phi_1)$. The theorem then says that $G(\Phi_1)$ has two kinds of real representation: canonical bias representations arising from $\mathbb{R}^*$-gains for $\langle\Phi\rangle$, in which the three edge lines are nonconcurrent, and canonical lift representations arising from $\mathbb{R}^+$-gains for Φ_1, in which the three edge lines concur. For this to actually occur there must be additive real gains for $\langle\Phi_1\rangle$, and indeed there are.

Example 7.2 (*Example 4.1 continued*). Let Φ_2 be Φ of Example 4.1 without h_1 and $(-1)e_{22}$. Most of Example 4.1 remains valid, except that there are noncanonical representations of $L(\Phi_2)$. This is due, as in the previous example, to the isomorphism of $L(\Phi_2)$ with $G(\Phi_2)$. Thus all canonical bias representations of Φ_2 are representations of $L(\Phi_2)$ but are not canonical lift representations because the edge lines do not concur. Still, by Theorem 7.1 there are no other representations but canonical lift and bias ones.

Example 7.3 (*Integral gains*). Take Ω to be $\langle \mathbb{Z}K_3 \rangle$ or any subgraph $\langle \Phi \rangle$ for which $\Phi \subseteq \mathbb{Z}K_3$ is large enough to force unique gains. (I do not know whether such a Φ can be finite; this is an interesting question for research.) Then any two nodes are multiply adjacent, so that $G(\Omega)$ has exactly two kinds of inequivalent representations (but not just two inequivalent representations) over $F = \mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$. The first kind is the canonical lift representations. All are projectively equivalent because they differ only by the embedding $\mathbb{Z} \hookrightarrow F$, which is determined by the image of 1; scaling the 0th coordinate is the equivalence. Second, there are the canonical bias representations, one for each possible image α of 1, namely any $\alpha \neq 0, \pm 1$, in the embedding $\mathbb{Z} \hookrightarrow F$. All are inequivalent, since scaling cannot transform one into another.

Example 7.4 (*Whirls*). Tutte's whirl matroid $\mathcal{W}_n = G(C_n^\bullet, \varnothing)$ is another which, by Theorem 7.1, has only canonical representations. This means its representations over any skew field are (up to projective equivalence) a 1-parameter family indexed by the orbits of $F^* \setminus \{1\}$ under the action of $\operatorname{Aut} F$. To see this we switch so that all edges but one in the C_n have gain 1, and let α be the gain of the remaining edge. Clearly, $\alpha \neq 1$ is the only restriction, and canonical bias representations corresponding to α and α' are projectively equivalent if and only if α and α' are equivalent by an automorphism (Proposition 2.9).

In particular, if $R(q)$ is the number of inequivalent representations over $\mathbb{F}_q$, we have $R(2) = 0$, $R(3) = 1$, and in general[7]

$$R(q) = \begin{cases} q-2 & \text{if } q \text{ is prime,} \\ -2 + \displaystyle\sum_{c|d} \frac{p^c}{c} \prod_{\substack{p' | \frac{d}{c} \\ \text{prime}}} \left(1 - \frac{1}{p'}\right) & \text{if } q = p^d. \end{cases}$$

Example 7.5 (*Signed expansion of K_3, continued from Example 2.4*). By Theorem 7.1 and the uniqueness of its gains (Theorem V.2.1), $\langle \pm K_3 \rangle$ has a canonical bias representation over F if and only if char $F \neq 2$ and a canonical lift representation if and only if char $F = 2$, and each of these is the unique F-representation. It also has representations over every field due to the isomorphism of $G(\pm K_3)$ with $G(K_4)$. If we examine the representation of $G(K_4)$ (which is unique since the matroid is unimodular [9]), we find that it is a canonical lift representation of $\pm K_3$ when char $F = 2$ and a canonical bias representation of $\pm K_3$ when char $F \neq 2$—just as it should be. See Figure 7.1.

Quite a different aspect of Theorem 7.1 is its conclusion that M can only be L or G—that is, there are no true intermediate matroids. This complements the similar Proposition II.4.4, which concerns biased graphs that are full and complete but assumes nothing about representability of M. In general, I believe, truly intermediate matroids cannot exist except on graphs of low connectivity, possibly 2 or less. On a separable graph they can exist.

Example 7.6 (*True intermediate matroids*). Suppose Ω has cutpoints $u_1, \ldots, u_k$ at which a central part, Ω_0, is separated from $\Omega_1, \ldots, \Omega_k$, respectively. Just to make the description easier, suppose also that each u_i supports a half edge h_i. Let M be the result of carrying out parallel connection of $G(\Omega_0)$ with $L(\Omega_i)$ at h_i for all $i > 0$. Then M is neither $G(\Omega)$

[7] I thank Dikran Karagueuzian and Marcin Mazur for calculating the number of orbits of $\mathbb{F}_q$.

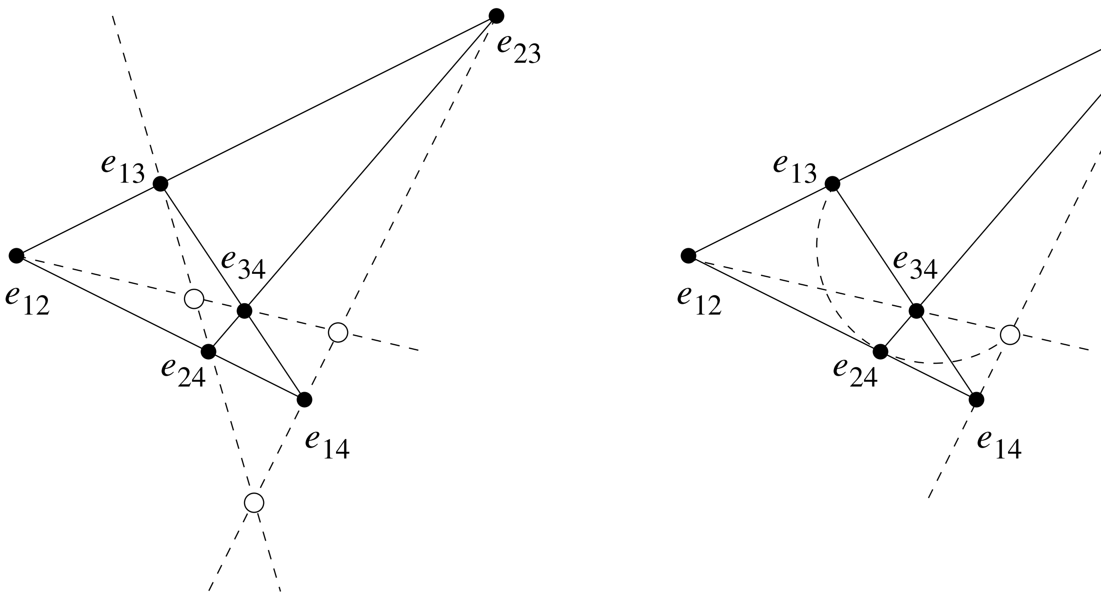

FIGURE 7.1. The two kinds of representation in a plane of $G(K_4)$, hence of $G(\pm K_3) = L(\pm K_3)$. They are, respectively, canonical bias and lift representations of $\pm K_3$.

nor $L(\Omega)$ except in special cases, but it is intermediate between them. Furthermore, M is representable over a skew field F if (and only if) all of $G(\Omega_0), L(\Omega_1), \ldots, L(\Omega_k)$ are.

8. Seven dwarves: Representations of the biased K_4's

We complete here the treatment of the seven biased graphs with underlying graph K_4, initiated in Part I and carried through Parts II and III. We call these biased graphs $\Omega_i = \Omega_i(K_4)$ for $i = 1, 2, \ldots, 7$. Recapitulating their definitions: $\Omega_1 = \langle K_4 \rangle$ is balanced; $\Omega_7 = (K_4, \varnothing)$ is contrabalanced. Ω_2 has two balanced triangles and consequently one balanced quadrilateral, the sum of the triangles. Ω_3 has the one balanced triangle $v_1v_2v_3$ and no balanced quadrilaterals. Ω_4, Ω_5, and Ω_6 have no balanced triangles and, respectively, three, two, and one balanced quadrilateral; then $\Omega_4 = \langle -K_4 \rangle$.

We wish to know the skew fields F over which each of these has canonical bias and lift representations as well as noncanonical matroid representations. The question is simplified by the fact that $G(\Omega_i) = L(\Omega_i)$, because K_4 contains no two node-disjoint circles. Therefore we have only one matroid to represent, but there are two distinct kinds of canonical representation that may or may not coincide. The facts are presented in Table 8.1.

Section III.13c presents the chromatic and balanced chromatic polynomials. For those matroids $G(\Omega_i)$ that have a real or complex canonical representation, these polynomials equal the characteristic polynomials of canonical hyperplanar representations. The unsigned coefficients of $\lambda^4, \lambda^3, \ldots, \lambda^0$ in the chromatic polynomials are the Betti numbers $\beta_0, \beta_1, \ldots, \beta_4$ of the complement of a complex hyperplanar bias representation as in Corollary 2.2. Similarly, the unsigned coefficients in the balanced chromatic polynomial are the Betti numbers of the complement of a complex affinographic representation $\mathcal{A}(\Omega_i)$ as in Corollary 4.5 (so that here $\beta_4 = 0$). As for real representations, $|\chi_{\Omega_i}(-1)|$ is the number of regions of a hyperplanar

	Ω_1	Ω_2	Ω_3	Ω_4	Ω_5	Ω_6	Ω_7
No can. bias rep.	–	$\mathbb{F}_2$	$\mathbb{F}_2, \mathbb{F}_3$	char $= 2$	$\mathbb{F}_2, \mathbb{F}_3$	$\mathbb{F}_2, \mathbb{F}_3, \mathbb{F}_4$	$\mathbb{F}_2, \mathbb{F}_3, \mathbb{F}_4$
No can. lift rep.	–	–	$\mathbb{F}_2$	char $\neq 2$	char $= 2$	$\mathbb{F}_2, \mathbb{F}_3$	$\mathbb{F}_2, \mathbb{F}_3, \mathbb{F}_4$
No rep.	–	–	$\mathbb{F}_2$	–	$\mathbb{F}_2$	$\mathbb{F}_2, \mathbb{F}_3$	$\mathbb{F}_2, \mathbb{F}_3, \mathbb{F}_4$

TABLE 8.1. The skew fields over which there do not exist canonical or any representations of the matroid $G(\Omega_i) = L(\Omega_i)$ for $\Omega_i = \Omega_i(K_4)$.

	Bias representation		Lift representation	
	Real	Complex	Real	Complex
	r	$\beta_0, \beta_1, \beta_2, \beta_3, \beta_4$	r	$\beta_0, \beta_1, \beta_2, \beta_3$
Ω_1	24	1, 6, 11, 6, 0	24	1, 6, 11, 6
Ω_2	38	1, 6, 13, 13, 5	30	1, 6, 13, 10
Ω_3	44	1, 6, 14, 16, 7	34	1, 6, 14, 13
Ω_4	46	1, 6, 15, 17, 7	(35)	(1, 6, 15, 13)
Ω_5	48	1, 6, 15, 18, 8	36	1, 6, 15, 14
Ω_6	50	1, 6, 15, 19, 9	37	1, 6, 15, 15
Ω_7	52	1, 6, 15, 20, 10	38	1, 6, 15, 16

TABLE 8.2. The number of regions (r) and the complementary Betti numbers (β_j) of a hyperplanar representation of Ω_i which is either a bias representation in $\mathbb{R}^4$ or $\mathbb{C}^4$ (left) or an affinographic representation in $\mathbb{A}^4(\mathbb{R})$ or $\mathbb{A}^4(\mathbb{C})$ (right). (An affinographic representation of Ω_4 does not exist in characteristic 0; the table shows the numbers that would apply if it did.)

representation in $\mathbb{R}^4$ and $|\chi^{\mathrm{b}}_{\Omega_i}(-1)|$ is that of an affinographic representation in $\mathbb{A}^4(\mathbb{R})$. All these numbers are in Table 8.2. Bear in mind that by a representation of Ω_i we mean a representation of any F^*- or F^+-gain graph Φ whose biased graph $\langle\Phi\rangle$ equals Ω_i. There may be many such gain graphs; see Section I.7.

Table 8.1 is justified by reasoning that has three parts: Table I.7.1 shows the possible gain groups of a gain graph for Ω_i, from which we immediately deduce the precise skew fields over which Ω_i has a canonical bias or lift representation; well-chosen contractions show the nonexistence of representations over certain small fields; and Proposition 8.1 assures us that there are no representations other than canonical ones.

The well-chosen contractions are

$$\begin{aligned} G(\Omega_3/e_{14}e_{24}) &= L_4, \\ G(\Omega_5/e_{13}e_{23}) &= L_4, \\ G(\Omega_6/e_{13}) = U_{3,5} &= L_5^\perp, \\ G(\Omega_7) = U_{4,6} &= L_6^\perp, \end{aligned}$$

where L_k is a k-point line (a uniform matroid of rank two). Since representing L_k or $L_k^\perp$ requires that $|F| \geq k-1$, the last line of Table 8.1 follows at once.

Proposition 8.1. *For each $i = 1, 2, \ldots, 7$, every representation of Ω_i is a canonical lift representation or a canonical bias representation. When $i \leq 4$, every representation is both* *canonical bias and canonical lift, whenever both kinds exist (see Table 8.1). When $i \geq 5$, no representation is simultaneously canonical for bias and lift.*

Outline of proof. The lift matroid of Ω_i for $i = 1, 2, 4$ is binary: graphic for $i = 1, 2$ and a signed-graphic lift matroid for Ω_4. Binary matroids are uniquely representable up to projective equivalence over any skew field for which a representation exists [9]; therefore the canonical bias and lift representations of Ω_1, Ω_2, and Ω_4 are the same whenever both exist.

In the other cases there is a node at which every triangle is unbalanced. We assume a projective representation and use the notation of Section 7, in which $\hat{e}$ represents edge e.

Lemma 8.2. *In Ω_i, if every triangle incident to node v_j is unbalanced, then the intersection of the representations of these triangles is a single point.*

We omit the proof. The point we call p_j.

Now, if $i \geq 5$ every node has a point p_j. It is clear that either all the points coincide or they have rank 4. One can verify that in the former case the representation extends to one of $L_0(\Omega_i)$ in which the p_j represent the extra point e_0, while in the latter case it extends to one of $G(\Omega_i^\bullet)$ in which each p_j represents the half edge h_j at v_j.

In the case of Ω_3 Lemma 8.2 yields a point p_4. One can prove that, p_4 representing e_0, one has a representation of $L_0(\Omega_3)$; this is similar to the proof when $i \geq 5$. Thus any representation of Ω_3 is a canonical lift representation. To conclude the proof we must show that this is at the same time a canonical bias representation so long as $|F| \geq 4$. We let p_4 represent h_4. We need to locate points p_2 and p_3 representing h_2 and h_3; then a point p_1 representing h_1 will be easy to find. We know $p_2 \in \overline{p_4\hat{e}_{24}}$ and $p_3 \in \overline{p_4\hat{e}_{34}}$. However, certain points on those lines must be avoided. The exclusions can be translated into a list of five points on $\overline{p_4\hat{e}_{24}}$ that cannot be p_2. If $|F| \geq 5$, that suffices to show p_2 exists. If $F = \mathbb{F}_4 = \{0, 1, \omega, \omega^2\}$, we know the exact representation because, being a canonical lift representation, it arises from F^+-gains, and we know all possible $\mathbb{F}_4^+$-gains up to switching: they are $\varphi(e_{ij}) = 0$ except for $\varphi(e_{24}) = \omega^k$ and $\varphi(e_{34}) = \omega^{k+1}$ for some k. Scaling lets us take $k = 0$. Now, with a single concrete representation in hand, we can show that two of the five forbidden points on $\overline{p_4\hat{e}_{24}}$ coincide; that is, a possible p_2 exists. Thus, the representation of $G(\Omega_3)$ extends to one of $G(\Omega_3^\bullet)$. $\square$

Finally, an unanswered question. Are all the canonical bias representations of Ω_i actually different, that is, projectively inequivalent? For Ω_1, Ω_2, and Ω_4 they are not, and for the others they might not be although the gains are not unique up to switching (by Section I.7). Similar remarks apply to canonical lift representations. Indeed the Ω_i for $i = 3, 5, 6, 7$ may be counterexamples to Conjectures 2.8 and 4.8, but this is unexplored.

9. Questions

Certain questions naturally arise from our work, that refine the general problems of representability raised in Problems II.2.13–14 and II.3.13. The greatest lack in our theory, and a difficult problem, is the question of noncanonical representation of the bias and lift matroids. Since there are natural ways to represent these matroids—by canonical representations—the most natural question in the world is: which gain or biased graphs allow a noncanonical representation? Although we can answer this for some very special (though large) classes of biased graphs, namely full and thick graphs, in general we are helpless. Even regarding such an obvious example as biased complete graphs, whose large number of edges ought to permit some deductions restricting possible representations, we can say nothing.

Consider the distinction between canonical and noncanonical representations. Results like Propositions 2.4 and 4.3 and especially Theorem 7.1 show that many biased graphs have only those representations, derived from gain functions in a skew field, that we call 'canonical'. Yet still these are only special kinds of biased graph.

Problem 9.1 (*Problems 2.5 and 4.4*). Which biased graphs have noncanonical representations?

Since canonical representations depend on having suitable gains, and since representability over finite fields is generally interesting, we ask about gains (for bias representation) in the multiplicative group of a Galois field, more generally in cyclic groups, and (for lift representation) in its additive group, which is a vector space over a prime field.

Problem 9.2. In which cyclic groups does a given biased graph have gains? Which biased graphs have gains in a given cyclic group?

Problem 9.3. In which vector spaces $\mathbb{F}_p^d$ does a given biased graph have gains? Which biased graphs have gains in a given space $\mathbb{F}_p^d$?

These questions are part of our legacy to the reader.

Errata for Parts I and III

Part I, p. 48, bottom: "Example 7.i below."

Part III, p. 45: "substituting $w = 0$, $v = -1$, and $w, x, \lambda = 1$ in the appropriate places."

Part III, p. 78, Example 12.2: Γ must be a simple graph.

Part III, p. 79: In (12.2) and (12.3), $(wx + 1)^n$ should be $(wx + v)^n$. Also, m denotes $\#E$. In the specializations, $(wx + 1)^n$ should be $(wx - 1)^n$ and $(x + 1)^n$ should be $(x - 1)^n$.

Part III, p. 83: In (13.3), $q_i^{[\mathrm{b}]}$ should be $\Delta q_i^{[\mathrm{b}]}$.

References

[1] N. Altshiller-Court, "College Geometry: A Second Course in Plane Geometry for Colleges and Normal Schools", Johnson, Richmond, Va., 1925; second ed., "College Geometry: An Introduction to the Modern Geometry of the Triangle and the Circle," Barnes & Noble, New York, 1952. [MR 14,307d]

[2] C. A. Athanasiadis, Characteristic polynomials of subspace arrangements and finite fields, *Adv. Math.* **122** (1996), 193–233. [MR 97k:52012]

[3] R. Baer, "Linear Algebra and Projective Geometry," Academic Press, New York, 1952. [MR 14,675j]

[4] R. E. Bixby, l-matrices and a characterization of binary matroids. *Discrete Math.* **8** (1974), 139–145. [MR 49 #2430]

[5] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler, "Oriented Matroids" (Encyc. Math. Appl., Vol. 46), Cambridge Univ. Press, Cambridge, 1993. [MR 95e:52023] Second ed., 1999. [MR 2000j:52016]

[6] R. G. Bland and M. Las Vergnas, Orientability of matroids, *J. Combin. Theory Ser. B* **24** (1978), 94–123. [MR 58 #5294]

[7] P. Boldescu, The theorems of Menelaus and Ceva in an n-dimensional affine space (in Romanian; French summary), *An. Univ. Craiova Ser. a IV-a* **1** (1970), 101–106. [MR 48 #12251]

[8] T. Brylawski, Constructions, in N. White, ed., "Theory of Matroids" (Encyc. Math. Appl., Vol. 26), Ch. 7, pp. 127–223, Cambridge Univ. Press, Cambridge, 1986. [MR 87k:05054 (book)]

[9] T. H. Brylawski and D. Lucas, Uniquely representable combinatorial geometries, in "Colloquio Internazionale sulle Teorie Combinatorie" (Proc., Rome, 1973), Vol. I, pp. 83–104, Atti dei Convegni Lincei, 17, Accademia Nazionale dei Lincei, Rome, 1976. [MR 55 #113 (book)]

[10] B. Budinský and Z. Nádeník, Mehrdimensionales Analogon zu Sätzen von Menelaos und Ceva, *Časopis Pěst. Mat.* **97** (1972), 75–77, 95. [MR 46 #6141]

[11] H.S.M. Coxeter, "Introduction to Geometry," second ed., Wiley, New York, 1969. [MR 49 #11369]

[12] H. H. Crapo and G.-C. Rota, "On the Foundations of Combinatorial Theory: Combinatorial Geometries," prelim. ed., MIT Press, Cambridge, Mass., 1970. [MR 45 #74]

[12a] G. Ding, B. Oporowski, J. Oxley, and D. Vertigan, Unavoidable minors of large 3-connected binary matroids, *J. Combin. Theory Ser. B* **66** (1996), 334–360 [MR 96m:05043].

[12b] G. Ding, B. Oporowski, J. Oxley, and D. Vertigan, Unavoidable minors of large 3-connected matroids, *J. Combin. Theory Ser. B* **71** (1997), 244–293 [MR 98h:05052].

[13] T. A. Dowling, A class of geometric lattices based on finite groups, *J. Combin. Theory Ser. B* **14** (1973), 61–86. [MR 46 #7066. Erratum, *ibid.* **15** (1973), 211. [MR 47 #8369]

[14] T. A. Dowling and D. G. Kelly, Elementary strong maps between combinatorial geometries, in "Colloquio Internazionale sulle Teorie Combinatorie" (Proc., Rome, 1973), Vol. II, pp. 121–152, Atti dei Convegni Lincei, 17, Accademia Nazionale dei Lincei, Rome, 1976. [MR 58 #27567]

[15] J. Geelen, A. M. H. Gerards, and G. Whittle, Branch-width and well-quasi-ordering in matroids and graphs, *J. Combin. Theory Ser. B* **84** (2002), 270–290. [MR 2003f:05027]

[16] A. M. H. Gerards, "Graphs and Polyhedra: Binary Spaces and Cutting Plane," CWI Tract 73, Centrum voor Wiskunde en Informatica, Amsterdam, 1990. [MR 92f:52027]

[17] C. Greene, Acyclic orientations (Notes), in M. Aigner, ed., "Higher Combinatorics" (Proc. NATO Adv. Study Inst., Berlin, 1976), pp. 65–68, NATO Adv. Study Inst. Ser., Ser. C, Vol. 31, D. Reidel, Dordrecht, 1977.

[18] C. Greene and T. Zaslavsky, On the interpretation of Whitney numbers through arrangements of hyperplanes, zonotopes, non-Radon partitions, and orientations of graphs, *Trans. Amer. Math. Soc.* **280** (1983), 97–126. [MR 84k:05032]

[19] F. Harary, On the notion of balance of a signed graph, *Michigan Math. J.* **2** (1953–1954), 143–146. Addendum, *ibid.*, preceding p. 1. [MR 16, 733]

[20] K. Hässig, "Graphentheoretische Methoden des Operations Research," B.G. Teubner, Stuttgart, 1979. [MR 80f:90002]

[21] J. P. Hogendijk, Mathematics in medieval Islamic Spain, in "Proceedings of the International Congress of Mathematicians" (Zürich, 1994), Vol. II, pp. 1568–1580, Birkhäuser, Basel, 1995. [MR 97c:01010]

[22] D. R. Hughes and F. C. Piper, "Projective Planes," Graduate Texts in Mathematics, Vol. 6l Springer-Verlag, New York-Berlin, 1973. [MR 48 #12278]

[23] T. Klein, A certain generalization of the theorems of Menelaos and Ceva (in Slovak), *Časopis Pěst. Mat.* **98** (1973), 22–25. [MR 47 #5704]

[24] E. G. Kogbetliantz, "Fundamentals of mathematics from an advanced viewpoint," 4 Vols. in 2, Gordon and Breach, New York, 1969.

[25] M. Las Vergnas, Sur les activités des orientations d'une géométrie combinatoire, in "Colloque Math. Discrètes: Codes et Hypergraphes" (Actes, Bruxelles, 1978), *Cahiers Centre Études Rech. Opér.* **20** (1978), 293–300. [MR 81d:05019]

[26] M. Las Vergnas, Convexity in oriented matroids, *J. Combin. Theory Ser. B* **29** (1980), 231–243. [MR 82f:05027]

[27] E. Minieka, "Optimization Algorithms for Networks and Graphs," Marcel Dekker, New York, 1978. [MR 80a:90066]

[28] P. Orlik and L. Solomon, Combinatorics and topology of complements of hyperplanes, *Invent. Math.* **56** (1980), 167–189. [MR 81e:32015]

[29] P. Orlik and H. Terao, "Arrangements of hyperplanes," Grundl. math. Wiss., Vol. 300, Springer-Verlag, Berlin, 1992. [MR 94e:52014]

[30] J. Oxley, D. Vertigan, and G. Whittle, On inequivalent representations of matroids over finite fields, *J. Combin. Theory Ser. B* **67** (1996), 325–343. [MR 97d:05052]

[30a] J. Oxley, D. Vertigan, and G. Whittle, On maximum-sized near-regular and $\sqrt[6]{1}$-matroids, *Graphs Combin.* **14** (1998), 163–179 [MR 99i:05058].

[31] D. C. Slilaty, "Orientations of Biased Graphs and Their Matroids," doctoral dissertation, Binghamton Univ., 2000.
[32] J. Sulzberger, personal communication, 1977.
[33] W. T. Tutte, Lectures on matroids, *J. Res. Nat. Bur. Standards (U.S.A.) Sect. B* **69B** (1965), 1–47. [MR 31 #4023]
[34] O. Veblen and J. W. Young, "Projective geometry," Vol. 1, second impression, Ginn, Boston, 1916; repr. by Blaisdell (Ginn), New York, 1965. [MR 31 #3912a]
[35] G. Whittle, talk at the Ohio State Univ. Conf. on Matroid Structure Theory, Columbus, O., 2002.
[36] T. Zaslavsky, "Facing up to arrangements: Face-count formulas for partitions of space by hyperplanes," Mem. Amer. Math. Soc. No. 154, American Mathematical Society, Providence, R.I., 1975. [MR 50 #9603]
[37] T. Zaslavsky, Signed graphs, *Discrete Appl. Math.* **4** (1982), 47–74. Erratum, *ibid.* **5** (1983), 248. [MR 84e:05095]
[38] T. Zaslavsky, Chromatic invariants of signed graphs, *Discrete Math.* **42** (1982), 287–312. [MR 84h:05050b]
[39] T. Zaslavsky, Bicircular geometry and the lattice of forests of a graph, *Quart. J. Math. Oxford* (2) **33** (1982), 493–511. [MR 84h:05050c]
[40] T. Zaslavsky, The biased graphs whose matroids are binary. *J. Combin. Theory Ser. B* **42** (1987), 337–347. [MR 88h:05082]
[41] T. Zaslavsky, Biased graphs. I. Bias, balance, and gains. II. The three matroids. III. Chromatic and dichromatic invariants, *J. Combin. Theory Ser. B* **47** (1989), 32–52; **51** (1991), 46–72; **64** (1995), 17–88. [MR 90k:05138; 91m:05056; 96g:05139]
[42] T. Zaslavsky, Biased graphs whose matroids are special binary matroids, *Graphs Combin.* **6** (1990), 77–93. [MR 91f:05097]
[43] T. Zaslavsky, Orientation of signed graphs, *European J. Combin.* **12** (1991), 361–375. [MR 93a:05065]
[44] T. Zaslavsky, Frame matroids and biased graphs, *European J. Combin.* **15** (1994), 303–307. [MR 95a:05021]
[45] T. Zaslavsky, A mathematical bibliography of signed and gain graphs and allied areas, *Electronic J. Combin.* (1998), Dynamic Surveys in Combinatorics, No. DS8 (electronic). [MR 2000m:05001a]
[46] T. Zaslavsky, Supersolvable frame-matroid and graphic-lift lattices, *European J. Combin.* **22** (2001), 119–133. [MR 20001k:05051]
[47] T. Zaslavsky, Perpendicular dissections of space, *Discrete Comput. Geom.* **27** (2002), 303–351. [MR 2003i:52026]
[48] T. Zaslavsky, Biased graphs. V. Group and biased expansions. VI. Synthetic geometry. VII. Contrabalance and antivoltages. VIII. A cornucopia of examples, in preparation.

Binghamton University, Binghamton, New York 13902-6000