BIASED GRAPHS. VII.
CONTRABALANCE AND ANTIVOLTAGES

THOMAS ZASLAVSKY

ABSTRACT. We develop linear representation theory for bicircular matroids, a chief example being a matroid associated with forests of a graph, and bicircular lift matroids, a chief example being a matroid associated with spanning forests. (These are bias and lift matroids of contrabalanced biased graphs.) The theory is expressed largely in terms of *antivoltages* (edge labellings that defy Kirchhoff's voltage law) with values in the multiplicative or additive group of the scalar field. We emphasize antivoltages with values in cyclic groups and finite vector spaces since they are crucial for representing the matroids over finite fields; and integer-valued antivoltages with bounded breadth since they are crucial in constructions. We find bounds for the existence of antivoltages and we solve some examples. Other results: The number of antivoltages in an abelian group is a polynomial function of the group order, and the number of integral antivoltages with bounded breadth is a polynomial in the breadth bound. We conclude with an application to complex representation. There are many open questions.

CONTENTS

Date: March 2001; Aug., 2002; Aug.–Sept., 2004; Nov.–Dec., 2005. Version of April 3, 2007.

2000 *Mathematics Subject Classification. Primary* 05B35, 05C22; *Secondary* 05A15, 05C35.

Key words and phrases. Antivoltage, bicircular matroid, bicircular lift matroid, forest matroid, spanning forest matroid, biased graph, gain graph, contrabalance, bias matroid, frame matroid, matroid representation, transversal matroid.

Research assisted by grant DMS-0070729 from the National Science Foundation.

Acknowledgement. I thank a referee of Part IV for pointing out the need for numerous illustrative examples, and the wonderful referee of this part who prepared two detailed lists of improvements based on notably thorough readings.

INTRODUCTION

This article concerns the problem of vector representation of two matroids associated with the forests of a graph and of their submatroids.

Consider a graph $\Gamma = (N, E)$, with node set N and edge set E. The *lattice of forests*, $\mathcal{F}(\Gamma)$, is the set of all forests in Γ, partially ordered by

$$F_1 \leq F_2 \iff N(F_2) \subseteq N(F_1) \text{ and every tree of } F_1 \text{ is contained in a tree of } F_2 \text{ or is node-disjoint from } F_2.$$

(Thus the top element of the forest lattice is the empty graph, $\varnothing$; the bottom element is the edgeless spanning forest $(N, \varnothing)$.) The *spanning forest lattice*, $\mathcal{F}_0(\Gamma)$, has two parts: the set $\mathcal{F}_{\mathrm{sp}}(\Gamma)$ of spanning forests and the set $\Pi(\Gamma)$ of all partitions of N into connected blocks; that is, for each $B \in \pi \in \Pi(\Gamma)$ the induced subgraph on B must be connected. The first part is ordered by inclusion of edge sets (since the node sets are the same), the second by refinement, and $F \leq \pi$ if and only if $\pi(F)$ refines π, where the blocks of $\pi(F)$ are the node sets of the components of F. (The top element of the spanning forest lattice is the partition 1_Γ of N induced by the components of Γ; the bottom element is $(N, \varnothing)$.) Both the forest lattice and the spanning forest lattice are geometric lattices; $\mathcal{F}_{\mathrm{sp}}(\Gamma)$ is a geometric semilattice that we might call the *spanning forest semilattice*. The forest lattice was introduced in [13]; it is the lattice of flats of a bicircular matroid which contains that of Γ as a submatroid and which we call the *forest matroid* of Γ, written (for reasons to be explained) $G(\Gamma^\circ, \varnothing)$. The spanning forest lattice comes from [14]; its matroid, which we call the *spanning forest matroid* of Γ and write as $L(\Gamma^\circ, \varnothing)$, is an analog of the bicircular matroid in which the circuits may be disconnected. These matroids and their submatroids are one of the principal subfamilies of the matroids of biased graphs (see Parts I–IV [15], of which however this article is largely independent). Our primary purpose is to investigate the implications for this family of the general theory of linear representations of biased-graphic matroids from Part IV. We are led thereby to a challenging new graph-theoretical problem: the analysis of antivoltages on a graph, by means of which we can investigate detailed properties of finite-field representations and derive bounds on their existence.

An *antivoltage* is a function $\varphi : E(\Gamma) \to \mathfrak{G}$ from oriented edges to a group, such that Kirchhoff's voltage law is everywhere violated; i.e., if $C = e_1 e_2 \cdots e_l$ is any circle in Γ with its edges in the indicated cyclic order, then

$$\varphi(e_1) + \varphi(e_2) + \cdots + \varphi(e_l) \neq 0,$$

the group identity. By saying "oriented edges" we mean that, if the orientation of edge e is reversed, the antivoltage is negated. We write the group additively because our main focus is on the multiplicative and additive groups of a finite field $\mathbb{F}_q$ of q elements. It is well known that $\mathbb{F}_q^*$, the multiplicative group of $\mathbb{F}_q$, is isomorphic to $\mathbb{Z}_{q-1}$ and that $\mathbb{F}_q^+$, the additive group, is isomorphic to $\mathbb{Z}_p^k$, where $q = p^k$. Theorems from Part IV imply that $\mathcal{F}(\Gamma)$ has a representation over $\mathbb{F}_q$ if and only if Γ has antivoltages in $\mathbb{F}_q^*$, and $\mathcal{F}_0(\Gamma)$ has such a representation if and only if Γ has antivoltages in $\mathbb{F}_q^+$. Thus antivoltages combinatorialize the representation problem for the forest and spanning forest matroids, and partially do so for bicircular and bicircular lift matroids in general.

A brief outline: In Section 1 is all necessary background from Parts I–IV. In the succeeding sections we develop increasingly less elementary lower and upper bounds on the size of fields over which the forest and spanning forest matroids are representable. The bound whose proof requires the most effort (Theorem 5.3) is that, for a simple graph of order n, $\mathcal{F}(\Gamma)$ has a representation over $\mathbb{F}_q$ for all $q \geq \lceil (e-2)(n-2)! \rceil$, where e is Euler's constant. (Still, there is ample room for improvement.) Section 4 shows that the best bounds—the critical parameters defined at the beginning of Section 3—respect the minor ordering of graphs. In Section 6 we do some examples that can be solved more completely, such as multilinks and wheels. Section 7 develops a theory of counting antivoltages, in which we show that the number of antivoltages with values in an abelian group is a polynomial function of the order of the group, while the number of integral antivoltages that satisfy a bound on the maximum total antivoltage around any circle is a polynomial function of the bound. Most sections raise unsolved problems.

There is another enumerative aspect to representations: One can count regions and faces of real hyperplanar representations of the bicircular or bicircular lift matroid, or find the Poincaré polynomial of the complement of a complex hyperplane representation. The main work there, according to Corollaries IV.2.3 and IV.4.5 (i.e., Corollaries 2.3 and 4.5 of Part IV), is the evaluation of the chromatic polynomial of $(\Gamma, \varnothing)$. For that we refer the reader to Example III.3.4 or, especially, [13].

1. Background

We begin by setting up notation for graphs. The notation $e{:}vw$ means that e is an edge with endpoints v and w and, if oriented, is oriented from v to w. A *link* has distinct endpoints and a *loop* has coinciding endpoints. The number of components of Γ is $c(\Gamma)$. The cyclomatic number is $\xi(\Gamma) = \#E - \#N + c(\Gamma)$. A *bond* is a minimal edge cutset. A graph is a *block* or *inseparable* if every pair of edges belongs to a common circle; equivalently, if it is connected and either is a single loop (with its supporting node) or has no loops and no cut nodes. A *block of* Γ is a maximal block subgraph of Γ.

The basis of this article, though mostly hidden, is the theory of biased graphs. A *biased graph* $(\Gamma, \mathcal{B})$ is a graph $\Gamma = (N, E)$ along with a list $\mathcal{B}$ of circles, called *balanced circles*, which has to satisfy a certain axiom whose statement is not needed here (see, e.g., Part I) but which is satisfied by the empty list. (A graph may have loops and multiple edges. The half and loose edges of previous parts will not be needed. We shall assume Γ is finite of order $\#N = n > 0$. A *circle* is the edge set of a simple closed path.) If the list is empty we have a *contrabalanced* biased graph $(\Gamma, \varnothing)$. Associated with $(\Gamma, \varnothing)$ are certain matroids, examples of the bias[1] and lift matroid constructions of Part II, that resemble the ordinary polygon or cycle matroid but in which every circle is independent.

The bias-matroid construction gives the *bicircular matroid* $G(\Gamma, \varnothing)$, introduced by Simões-Pereira [11] for finite graphs and in an infinite version by Klee [4] and further studied in [12], [5], etc. The lift-matroid construction gives the *bicircular lift matroid* $L(\Gamma, \varnothing)$ and its inseparable companion the *bicircular complete lift matroid* $L_0(\Gamma, \varnothing)$. In the bicircular matroid the circuits are the minimal connected edge sets of cyclomatic number 2; in the bicircular lift matroid the circuits are the minimal edge sets of cyclomatic number 2, not necessarily connected. $L_0(\Gamma, \varnothing)$ is the bicircular lift matroid together with a so-called *extra point* that behaves like a graph loop. (In this matroid all graph loops behave the same, and

[1]The name *frame matroid* would be preferable in view of [16], but we adhere to the terminology of Parts I–IV.

they are not matroid loops.) If we indicate by Γ° the graph with a loop adjoined to every node, then $G(\Gamma^\circ, \varnothing)$ is the matroid whose lattice of flats is the forest lattice; and $L(\Gamma^\circ, \varnothing)$ is the matroid whose lattice of flats is the spanning forest lattice. In the latter, all the loops collapse into one atom, which can be identified with the extra point; so it is equivalent to the bicircular complete lift matroid (except for the empty graph, $\varnothing$; we assume $n > 0$ to avoid this example). Consequently, we shall write $L(\Gamma^\circ, \varnothing)$ rather than $L_0(\Gamma, \varnothing)$.

These matroids are matroids on the edge set of the graph. They give rise to the forest and spanning-forest lattices in the following way. The graph is Γ°. A forest F corresponds to the edge set $E(F) \cup E(\Gamma^\circ \setminus N(F))$. Thus, the partial ordering in the forest lattice derives from containment of edge sets in the bicircular matroid of Γ°. A partition $\pi \in \Pi(\Gamma)$ corresponds to the edge set $E{:}\pi$ that is the union of all edge subsets induced by node sets $B \in \pi$. (So, π is recoverable from $E{:}\pi$ as the class of node sets of its components.) Thus, the partial ordering in the spanning-forest lattice also derives from edge-set containment, but the edge sets are only partly the same as with the forest lattice.

Switching an antivoltage mapping $\varphi : E \to \mathfrak{G}$ (from Section I.5) means choosing a function $\eta : N \to \mathfrak{G}$ and replacing φ by φ^η defined by $\varphi^\eta(e) = -\eta(v)+\varphi(e)+\eta(w)$ (in additive notation for future use). The basic facts are that any antivoltage can be switched to be zero on any forest (from Lemma I.5.3) without giving to any circle the identity antivoltage (Lemma I.5.2) and, in the abelian case, without even changing the antivoltage of any circle.

Our interest is in the fields (and skew fields) F, and mainly the finite fields, over which the bias and lift matroids have vector representations, especially those over which there are "canonical" representations. We need the following terminology: A *bias representation* of $(\Gamma, \varnothing)$ is a vector representation of $G(\Gamma, \varnothing)$ and a *lift representation* is a vector representation of $L(\Gamma, \varnothing)$. The definition of canonical representations begins with an antivoltage function φ on Γ. (In the general terms of Parts I–IV, antivoltages are called *gains for* $(\Gamma, \varnothing)$ and canonical representations are defined in terms of gains.) If φ has values in F^*, a *canonical bias representation* of $(\Gamma, \varnothing)$ is the bias representation in F^n for which a link $e{:}v_iv_j$ represents as a vector $b_i - \varphi(e)b_j$ and a loop $e{:}v_iv_i$ represents as b_i, or any scaling of such a representation. (The set $\{b_i\}$ is the standard unit coordinate basis of F^n. *Scaling* means multiplying each vector independently by a nonzero scalar.) If φ has values in F^+, a *canonical lift representation* of $(\Gamma, \varnothing)$ is any scaling of the lift representation in F^{n+1} in which a link $e{:}v_iv_j$ represents in F^{n+1} as a vector $-b_i + b_j + \varphi(e)b_0$, where $\varphi(e) \in F^+$, and a loop as the vector b_0. By Propositions IV.2.4 and IV.4.3 we have:

Fundamental Lemma. *(a) A bias representation of* $(\Gamma, \varnothing)$ *in dimension* n *is a canonical bias representation if and only if it extends to a representation of* $G(\Gamma^\circ, \varnothing)$*. Thus all representations of* $G(\Gamma^\circ, \varnothing)$ *are canonical with respect to some choice of antivoltage in* F^*.

(b) A lift representation of $(\Gamma, \varnothing)$ *in dimension* $n+1$ *is a canonical bias representation if and only if it extends to a representation of* $L(\Gamma^\circ, \varnothing)$*. Thus all representations of* $L(\Gamma^\circ, \varnothing)$ *are canonical with respect to some choice of antivoltage in* F^+.

What this means is that the statements

(G1) $(\Gamma, \varnothing)$ has a canonical bias representation over F,
(G2) $G(\Gamma^\circ, \varnothing)$ has a representation over F,
(G3) Γ has antivoltages (i.e., $(\Gamma, \varnothing)$ has gains) in F^*

are equivalent, and so are the statements

(L1) $(\Gamma, \varnothing)$ has a canonical lift representation over F,

(L2) $L(\Gamma^\circ, \varnothing)$ has a representation over F,
(L3) Γ has antivoltages (i.e., $(\Gamma, \varnothing)$ has gains) in F^+.

The problem of determining which skew fields admit canonical representations of $G(\Gamma, \varnothing)$ and $L(\Gamma, \varnothing)$ thereby becomes that of deciding which multiplicative or additive groups of skew fields admit antivoltages for Γ. Indeed, one can count canonical representations by counting antivoltages, because the different canonical representations (with respect to a fixed basis) correspond to the different antivoltages. One may even speculate that, for most graphs, projectively equivalent representations correspond to equivalence classes of antivoltages under switching and field automorphisms—see Conjectures IV.2.8 and IV.4.8. For the most part we shall not further mention the representational consequences of antivoltages, as they are implicit in the Fundamental Lemma.

For existence and (in part) enumeration of antivoltages it is enough to treat graphs that are inseparable and edge 3-connected. (The same holds for gains on any biased graph.) Regarding separability, an antivoltage on Γ is equivalent to having an antivoltage on each of the blocks of Γ. For edge 3-connection, we define a *series class* of edges as an equivalence class under the relation of belonging to a common two-edge bond. Let R be the complement of a system of distinct representatives for the family of series classes of Γ. (We assume for simplicity that Γ is inseparable.) Then the contraction Γ/R is edge 3-connected if it has more than one edge. Also, there is a natural bijection between the spanning trees of Γ/R and those of Γ that contain R. Since every antivoltage switches to be zero on R itself, antivoltages exist on Γ if and only if they exist on Γ/R. The number of antivoltages on Γ is determined by the number on Γ/R if the group is abelian, due to properties of switching that we discuss in Section 7. As the cyclomatic number seems to have an important role, we mention further that $\xi(\Gamma) = \xi(\Gamma/R)$.

Canonical representations are obviously important, because they are the only representations of the forest and spanning forest matroids. But if a canonical representation does not exist one will naturally ask about noncanonical ones, not derived from antivoltages. Except for the simplest lower bounds on the order of the field, that is too difficult for us to treat here. It seems likely that most well connected graphs will have only canonical representations; but for the present that is merely a conjecture.

2. General bounds on representability and antivoltages

Lower bounds on the sizes of fields over which bias and lift matroids have representations follow from restrictions imposed on antivoltages, or directly on representations, by biased-graph minors whose matroids are lines (uniform matroids of rank 2). The restrictions apply to all biased graphs but are especially applicable to contrabalanced ones. By mK_2 we mean an m-fold multiple link. By $\Omega^{(l)}$ we mean Ω with unbalanced loops attached to l nodes. The first lemma treats canonical representations; the second concerns any representation. We state the lemmas for biased graphs in general; for the present article one can think of Ω as simply $(\Gamma, \varnothing)$.

We restate the definitions of contraction and minors from Part I in a simplified form appropriate to our work here. To contract a link e in a biased graph Ω we identify its endpoints to a single node and delete the edge. A balanced circle of the contraction is a circle C that is a balanced circle in Ω or is a path in Ω that joins the endpoints of e and makes, with e, a balanced circle in Ω. To contract a loop e, one deletes it and its supporting node v and changes every link incident with v into a loop supported by its second endpoint; any other loops at v should be deleted, but note that links incident with v are not deleted. A finite edge set S can be contracted by contracting its edges one at a time in any order. We denote by Ω/S the result of contracting Ω by S. A *minor* is a subgraph of a contraction. Evidently, a minor of a contrabalanced graph is still contrabalanced.

Lemma 2.1. *Suppose a biased graph Ω contains $(mK_2, \varnothing)$ as a minor. Then it has no gains in any group of order less than m. Furthermore, Ω has no canonical bias representation over $\mathbb{F}_q$ for $q \leq m$ and no canonical lift representation for $q < m$.*

Proof. Clearly, gains for $(mK_2, \varnothing)$ must be in a group of order at least m. Gains for Ω transform to gains for its minors (Corollary I.5.7). Thus a gain group for Ω must have order m at least. □

Lemma 2.2. *Suppose a biased graph Ω has $(kK_2, \varnothing)$, $((k-1)K_2^{(1)}, \varnothing)$, or $((k-2)K_2^{(2)}, \varnothing)$ as a minor. Then $G(\Omega)$ is not representable over any field F with $\#F < k-1$. If Ω has either of the first two minors, then $L(\Omega)$ is not representable over any field F with $\#F < k-1$. If Ω has the first minor, then $L_0(\Omega)$ is not representable over any field F with $\#F < k$.*

Proof. The bias and lift matroids of $((k-l)K_2^{(l)}, \varnothing)$ equal the k-point line $U_{2,k}$ (except for the lift matroid when $l = 2$). Since Ω has $((k-l)K_2^{(l)}, \varnothing)$ as a minor, $G(\Omega)$ and $L(\Omega)$ have $U_{2,k}$ as a minor (by Theorems II.2.5 and II.3.6; for the lift case, we have to note that the minor can be obtained by contracting only a balanced set of edges).

Now the lemma is immediate from the well-known fact that for a matroid to be representable over $\mathbb{F}_q$ it cannot have any $U_{2,q+2}$ minor. □

Now we look for lower bounds on the size of antivoltage groups. An *augmented bond* in a graph is a bond together with an edge of one circle contained in each side of the bond that contains a circle. Thus if the bond separates two circles, the augmented bond is two edges larger than the bond. If deleting the bond leaves no circles, the augmented bond is the bond itself. From augmented bonds we get a lower bound on all representability, not necessarily canonical. A *semiaugmented bond* is a bond together with one edge of a circle contained in one side of the bond, if such a circle exists, or just the bond if there is no such circle. In the next two propositions, $M(\Gamma)$, $M'(\Gamma)$, and $M''(\Gamma)$ are the maximum sizes of a bond, a semiaugmented bond, and an augmented bond in Γ.

Lemma 2.3. *A graph Γ has antivoltages in $\mathfrak{G}$ only if $\#\mathfrak{G} \geq M(\Gamma)$. It has antivoltages in F^* only if $\#F > M(\Gamma)$ and in F^+ only if $\#F \geq M(\Gamma)$.*

Proof. Starting from $(\Gamma, \varnothing)$, one gets an $(mK_2, \varnothing)$ minor with $m = M(\Gamma)$ by contracting the complement of a largest bond. Apply Lemma 2.1. □

Lemma 2.4. *The bicircular matroid $G(\Gamma, \varnothing)$ is representable over F only if $\#F \geq M''(\Gamma)-1$. The bicircular lift matroid $L(\Gamma, \varnothing)$ is representable over F only if $\#F \geq M'(\Gamma)-1$, and $L(\Gamma^\circ, \varnothing)$ is representable over F only if $\#F \geq M(\Gamma)$.*

Proof. Apply Lemma 2.2. □

We should mention a weak but completely general upper bound on the order of a group in which Γ does *not* have antivoltages.

Proposition 2.5. *If ξ is the cyclomatic number of Γ, then there exist antivoltages for Γ in any group of order $\mu \geq 2^{\xi-1}(\xi-1)!\sqrt{e}$.*

Proof. Gagola [3] proved that any group $\mathfrak{G}$ of order at least $2^{\xi-1}(\xi-1)!\sqrt{e}$ contains ξ elements that avoid the identity element 1: no product of one or more of them or their inverses equals the identity, where no element can be repeated in a product, not even inverted. If we have such a set in $\mathfrak{G}$, we assign antivoltage 1 to a maximal forest and the ξ values in our set to the remaining ξ edges. □

3. Modular, integral, and prime-power antivoltages

For questions of representation certain groups are especially important. We call antivoltages *modular* if they have values in a finite cyclic group (μ-*modular* if the group is $\mathbb{Z}_\mu$) and *prime-power* if the values are in the additive group $\mathbb{Z}_p^k$ of $\mathbb{F}_{p^k}$. Some key numbers, which it is not difficult to see are well defined if Γ is not a forest (which we shall assume throughout), are

$\mu_0(\Gamma) =$ the smallest μ for which Γ has antivoltages in the cyclic group $\mathbb{Z}_\mu$,
$\mu_1(\Gamma) =$ the smallest μ for which Γ has antivoltages in every $\mathbb{Z}_{\mu'}$ with $\mu' \geq \mu$,
$\lambda_0(\Gamma) =$ the smallest prime power p^k such that Γ has antivoltages in $\mathbb{Z}_p^k$,
$\lambda_1(\Gamma) =$ the smallest prime power λ such that Γ has antivoltages in $\mathbb{Z}_p^k$ for every prime power $p^k \geq \lambda$,
$\kappa_p(\Gamma) =$ the smallest k such that Γ has antivoltages in $\mathbb{Z}_p^k$, where p is a prime number.

Note that Γ has antivoltages in $\mathbb{Z}_p$ for every $p \geq \mu_1(\Gamma)$ but no $\mathbb{Z}_p$ with $p < \mu_0(\Gamma)$. All our upper bounds on modular antivoltages are actually based on still another parameter. *Integral antivoltages* take values in $\mathbb{Z}$. Define

$\mu_2(\Gamma) =$ the minimum breadth of integral antivoltages on Γ,

the *breadth* of an integral antivoltage φ being defined as $1 + \max_C \varphi(C)$, where C ranges over all oriented circles in Γ. (If C has antivoltage α in one orientation, its antivoltage is $-\alpha$ in the opposite orientation; it follows that $\mu_2 > 0$.) If we switch φ to be zero on a maximal forest, then $|\varphi(e)| < \text{breadth}(\varphi)$ for all edges.

An antivoltage mapping modulo μ can be taken as integral with values in the interval $[-(\mu-1), \mu-1]$. Thus, any $\mathbb{Z}_\mu$-antivoltage φ_μ derives from an integral antivoltage φ by reducing the values modulo μ, although the breadth of φ may be larger than μ.

If Γ is not a block graph, then it has antivoltages in a group if and only if each of its blocks does. Therefore, each of the parameters $\mu_i, \lambda_i, \kappa_p$ on Γ equals the maximum of its values on the blocks of Γ. Thus in looking for bounds we can confine our attention to block graphs. Recall the maximum bond size $M(\Gamma)$ from Section 2.

Proposition 3.1. *We have*
(a) $\mu_0(\Gamma) \geq M(\Gamma)$,
(b) $\kappa_p(\Gamma) \geq \log_p M(\Gamma)$,
(c) $\lambda_0(\Gamma) \geq M(\Gamma)$.

Proof. By Lemma 2.3. □

The next result gives upper bounds whose virtue is in their simplicity. Generally speaking, and for μ_2 always, Theorem 3.4 is better. Still, these bounds are better than the ones derived from Proposition 2.5, except for $\xi \leq 4$ in (d).

Theorem 3.2. *If Γ is a graph whose cyclomatic number is ξ, then*

(a) $\mu_0(\Gamma) \leq \mu_1(\Gamma) \leq \mu_2(\Gamma)$,
(b) $\mu_2(\Gamma) \leq 2^\xi$,
(c) $\kappa_p(\Gamma) \leq \lceil \xi / \lfloor \log_2 p \rfloor \rceil$,
(d) $\lambda_1(\Gamma) \leq 3^\xi$.

Proof. In each case take the antivoltage equal to zero on a maximal forest. Define $l = \lfloor \log_2 p \rfloor$.

(a) An integral antivoltage φ of breadth k gives a modular antivoltage modulo any $\mu \geq k$.

(b) Assign integers $1, 2, 4, \ldots, 2^{\xi-1}$ to the remaining ξ edges. No combination of these nonzero numbers by addition and subtraction can equal 0 or be as large as 2^ξ. Therefore, no circle can have antivoltage equal to 0 or larger than $2^\xi - 1$.

(c) Let $b_1, \ldots, b_k$ be generators for $\mathbb{Z}_p^k$ and take the elements $2^j b_i$ for $0 \leq j \leq l-1$. No sum of one or more of these kl elements or their negatives equals zero. Thus, as long as $kl \geq \xi$, we can assign ξ values to the edges outside the forest and have an antivoltage.

(d) We need to show that an antivoltage exists in $\mathbb{Z}_p^k$ if $p^k \geq 3^\xi$. Since $3^{kl} \geq (2^{l+1}-1)^k \geq p^k$, the hypothesis implies that $kl \geq \xi$, so $k \geq \kappa_p$ and the antivoltage exists by (c). □

A *transversal matroid* is defined by partial transversals of a family N of subsets of a ground set E (see [7, Section 1.6]). If each triple $\{X, Y, Z\}$ of sets has empty intersection $X \cap Y \cap Z$, then $\Gamma = (N, E)$ defines a graph in which node-edge incidence is reverse set membership. Then $M = G(\Gamma, \varnothing)$, as Matthews observed [5].

Corollary 3.3. *If M is a transversal matroid on m points defined by a family of sets of which each three have empty intersection, then M is representable over every field of order at least 2^m.*

Proof. Note that $\xi(\Gamma) \leq m$. □

For transversal matroids of the specified type the corollary improves on the theorem of Piff and Welsh, who proved that each transversal matroid has a representation over every sufficiently large field, but without giving a specific bound [8].

A construction by degrees produces generally better but more complicated bounds. We may as well suppose Γ is a block of order $n \geq 2$. (We exclude loops because, if Γ does have loops, we get better bounds by removing them.) Choose a spanning tree T and an acyclic orientation α of $\Gamma \setminus T$. Let d_i^o be the outdegree of v_i and define

$$\pi(\Gamma \setminus T, \alpha) = \prod_{i=1}^{n} (d_i^o + 1),$$

$$\sigma_p(\Gamma \setminus T, \alpha) = \sum_{i=1}^{n} \lceil \log_p (d_i^o + 1) \rceil.$$

Theorem 3.4. *For Γ, T, and α as described, we have*

(a) $\mu_2(\Gamma) \leq \pi(\Gamma \setminus T, \alpha)$,
(b) $\kappa_p(\Gamma) \leq \sigma_p(\Gamma \setminus T, \alpha)$,
(b′) $\kappa_p(\Gamma) = 1$ *if* $p \geq \pi(\Gamma \setminus T, \alpha)$,
(c) $\lambda_1(\Gamma) < \pi(\Gamma \setminus T, \alpha)^2$.

Proof. Index the nodes $v_1, \ldots, v_n$ in order of increasing outdegree. (This is for the proof of part (c).) Set $\pi_j = \prod_{i=1}^{j} (d_i^o + 1)$; note that $\pi_0 = \pi_1 = 1$ and $\pi_n = \pi(\Gamma \setminus T, \alpha)$. Assign 0 to every edge in T.

(a) To get an antivoltage in $\mathbb{Z}_\mu$ we need values on $\Gamma\backslash T$ so that no circle in Γ has antivoltage $\varphi(C) = 0$. We do this by constructing an integral antivoltage of sufficient breadth. We assign to each edge departing v_i a different one of the integers $1\pi_{i-1}, 2\pi_{i-1}, \ldots, d_i^o\pi_{i-1}$. Each node v_i at which C contains a departing edge contributes $\pm k\pi_{i-1}$ to $\varphi(C)$, where $1 \leq k \leq d_i^o$. If amongst the nodes in C having an outgoing edge the one with the largest index is v_j, the other terms in $\varphi(C)$ total less than π_{j-1} in absolute value, while v_j contributes at least this much. Therefore, $\varphi(C) \neq 0$.

The largest conceivable value of $\varphi(C)$ is $d_2^o\pi_1 + d_3^o\pi_2 + \cdots + d_n^o\pi_{n-1} = \pi_n - 1$. Thus, φ has breadth at most π_n, proving (a).

(b) We adapt the idea of (a) to modular arithmetic. We assign to each v_i a group $\mathbb{Z}_p^{k_i}$. We define a function $\varphi_i : E \setminus T \to \mathbb{Z}_p^{k_i}$ by assigning d_i^o distinct nonzero members of $\mathbb{Z}_p^{k_i}$ to the d_i^o outward edges from v_i and zeroes to the other edges. This requires dimension $k_i = \lceil \log_p(d_i^o+1) \rceil$. Then $\varphi_1 \oplus \cdots \oplus \varphi_n$ is an antivoltage with values in the group $\mathbb{Z}_p^{k_1} \oplus \cdots \oplus \mathbb{Z}_p^{k_n}$. Therefore, $\kappa_p \leq \sum_i k_i$.

(b′) For a prime number $p \geq \pi_n$ we apply the method of (a) to construct antivoltages in $\mathbb{Z}_p$.

Before we prove (c) we take advantage of (b′) to develop another formula for λ_1.

Lemma 3.5. *λ_1 is the smallest prime power $\lambda > \max_p p^{\kappa_p - 1}$.*

Proof. The maximum is well defined because part (b′) shows that there are only finitely many primes for which $p^{\kappa_p-1} > 1$. Clearly, λ_1 is greater than the maximum, and nothing otherwise prevents it from being as small as possible so long as it is a prime power. □

Now we can prove part (c). Suppose p is the prime that maximizes p^{κ_p-1}. From (b′) we know $p < \pi_n$. Since $p < \pi_n$, there are indices $1 = j_0 < j_1 < \cdots < j_k < j_{k+1} = n$ (for some $k \geq 0$) and exponents $r_1, \ldots, r_{k+1} > 0$ such that

$$\begin{aligned} p^{r_1-1} < \pi_{j_1-1} &\leq p^{r_1} < \pi_{j_1}, \\ p^{r_2-1} < \pi_{j_2-1}/\pi_{j_1-1} &\leq p^{r_2} < \pi_{j_2}/\pi_{j_1-1}, \\ &\cdots \\ p^{r_k-1} < \pi_{j_k-1}/\pi_{j_{k-1}-1} &\leq p^{r_k} < \pi_{j_k}/\pi_{j_{k-1}-1}, \\ p^{r_{k+1}-1} < \pi_n/\pi_{j_k-1} &\leq p^{r_{k+1}}; \end{aligned}$$

and in addition, $r_m > 1$ only if $d_{j_m}^o \geq p$; $j_{m+1} = j_m + 1$ if $r_m > 1$ (for $m \leq k$); and k is as small as possible given the other requirements. Let $r = r_1 + \cdots + r_{k+1}$.

For each $m = 1, \ldots, k+1$ we use the method of (a) or (b) to define $\varphi_m : E\backslash T \to \mathbb{Z}_p^{r_m}$ which assigns distinct nonzero values in $\mathbb{Z}_p^{r_m}$ to the outward edges from each v_i for $j_m \leq i < j_{m+1}$ and the value 0 to all other edges. When $r_m = 1$, then to the edges departing v_i we assign the numbers $1\pi_{i-1}/\pi_{j_m-1}, 2\pi_{i-1}/\pi_{j_m-1}, \ldots, d_i^o\pi_{i-1}/\pi_{j_m-1}$, all of which are less than p, and then we interpret them as values in $\mathbb{Z}_p$. But when $r_m > 1$, we assign $d_{j_m}^o$ distinct nonzero values in $\mathbb{Z}_p^{r_m}$ to the outward edges from v_{j_m}.

The direct sum $\varphi_1 \oplus \cdots \oplus \varphi_{k+1} : E \setminus T \to \mathbb{Z}_p^r$ is an antivoltage mapping with group of order p^r. Now,

$$p^r = \prod_{m=1}^{k+1} p^{r_m} < \pi_{j_1}(\pi_{j_2}/\pi_{j_1-1}) \cdots (\pi_{j_k}/\pi_{j_{k-1}-1})(\pi_n/\pi_{j_k-1})p$$

because $p^{r_{k+1}-1} < \pi_n/\pi_{j_k-1}$,

$$\begin{aligned} &= (d^o_{j_1}+1)(d^o_{j_2}+1)\cdots(d^o_{j_k}+1)\pi_n p \\ &< \pi_n \cdot (d^o_{j_1}+1)(d^o_{j_2}+1)\cdots(d^o_{j_k}+1)\pi_{j_1} \end{aligned}$$

because $p < \pi_{j_1}$,

$$\begin{aligned} &\le \pi_n \cdot \pi_{j_1}(d^o_{j_1+1}+1)(d^o_{j_2+1}+1)\cdots(d^o_{j_k+1}+1) \\ &\le \pi_n \cdot \pi_{j_k+1} \le \pi_n^2. \end{aligned}$$

It follows that there is a power of p, $p^r < \pi(\Gamma \setminus T, \alpha)^2$, for which there exist antivoltages in $\mathbb{Z}_p^r$. This implies (c). □

The question arises of how to choose T and α most wisely. Richard Stanley (personal communication) suggested a partial answer. Think of α as an acyclic orientation of $\Gamma' = \Gamma \backslash T$.

Proposition 3.6. *If an orientation α of a finite graph Γ' minimizes $\pi(\Gamma', \alpha)$, then α is maximal in the dominance ordering of orientations.*

The *dominance ordering* $\succcurlyeq$ is a partial ordering of unordered partitions of m into n non-negative integers. Given two such partitions, d and d', we first arrange each in descending order and then define $d' \succcurlyeq d \iff d'_1 + \cdots + d'_i \geq d_1 + \cdots + d_i$ for every $i = 1, 2, \ldots, n$. We apply this to orientations through their outdegree sequences, getting a partial quasiordering since different orientations may have the same outdegrees. Define $\pi(d) = \prod_i (d_i + 1)$. Then

$$d' \succ d \implies \pi(d') < \pi(d). \tag{3.1}$$

To prove this, note that there are a smallest k such that $d'_1 + \cdots + d'_k > d_1 + \cdots + d_k$ and a smallest $l > k$ such that $d'_1 + \cdots + d'_l = d_1 + \cdots + d_l$. Then $d'_k > d_k \geq d_l > d'_l$. Set $d''_i = d'_i$ except for $d''_k = d'_k - 1$ and $d''_l = d'_l + 1$; then $\pi(d'') > \pi(d')$. The formula follows by induction on the height of the interval $[d, d']$ in the poset.

Proof of Proposition. Apply (3.1) to the outdegree sequences of orientations of Γ'. □

It is well known that all maximal orientations are acyclic; but not all acyclic orientations are maximal, and not every maximal orientation gives the minimum value of $\pi(\Gamma', \alpha)$. Thus, the problem of choosing the best α is not completely solved, nor is that of choosing the best T in Theorem 3.4.

Example 3.1 (*Multiple links*). Comparing μ_0 to μ_1, one readily sees that they are equal for the m-fold link mK_2 (or a subdivision; this includes circles other than loops, and theta graphs). For mK_2, Lemma 2.1 implies that $\mu_0 \geq m$. An obvious integral antivoltage assigns $0, 1, \ldots, m-1$ to the m edges. Hence,

$$\mu_0(mK_2) = \mu_1(mK_2) = \mu_2(mK_2) = m.$$

By similar reasoning,

$$\begin{aligned} \lambda_1(mK_2) &= \text{ the least prime power } q \geq m, \\ \kappa_p(mK_2) &= \lceil \log_p m \rceil. \end{aligned}$$

Thus, $(mK_2, \varnothing)$ has a canonical bias representation over $\mathbb{F}_q$ if and only if $q > m$, and a bias representation if and only if $q \geq m-1$. It has a canonical lift representation if and only if $q \geq m$, and a lift representation if and only if $q \geq m-1$. For the proof, observe that

$G(mK_2, \varnothing) = L(mK_2, \varnothing) = U_{2,m}$, the m-point line. The conclusions follow from this and Lemmas 2.1 and 2.2.

We also know that μ_0, μ_1, and μ_2 are equal for K_4, by Equation (6.2).

Problem 3.7. Are μ_0 and μ_1 equal for every finite graph? In other words, in the set of orders of cyclic groups in which Γ has antivoltages, are there no gaps?

I do not know whether any of the parameters are independent.

Problem 3.8. Is $\mu_2 = \mu_1$? A stronger question: If φ_μ is an antivoltage in $\mathbb{Z}_\mu$, does there exist an integral antivoltage $\varphi \equiv \varphi_\mu \pmod{\mu}$ with breadth $\leq \mu$?

4. Minors and matroids

In the *minor ordering* of graphs or matroids, $A \leq B$ if A is a minor of B.

Proposition 4.1. *The functions μ_i, λ_i, and κ_p of a graph Γ are determined by the polygon matroid $G(\Gamma)$. Furthermore, they are weakly increasing with respect to the minor orderings of graphs and of matroids.*

Proof. The first part is a special case of Theorem IV.5.1.

It suffices to prove the second part for graphs. Weak increase is obvious for subgraphs. The proof for contractions relies on three facts from Part I. First, $\Gamma/S = (\Gamma/T) \setminus (S \setminus T)$ if T is a maximal forest in $S \subseteq E$. Therefore we need only contract by a forest T. Second, $(\Gamma, \varnothing)/T = (\Gamma/T, \varnothing)$ for a forest. (Equality fails if T is not a forest.) Finally, gains contract: if Ω has gains in $\mathfrak{G}$, so does every contraction. □

Proposition 4.2. *For each skew field F, the class of graphs Γ for which $(\Gamma, \varnothing)$ has each of the following kinds of representation is closed under taking minors.*

1. *A canonical bias representation.*
2. *A bias representation.*
3. *A canonical lift representation.*
4. *A lift representation.*

Proof. Again we use the fact that $\Gamma/S = (\Gamma/T) \setminus (S \setminus T)$ if T is a maximal forest in S. □

We conclude, by the famous theorem of Robertson and Seymour [9, 10], that each class of finite graphs mentioned in Proposition 4.2, as well as each class of finite graphs defined by the property that $f(\Gamma) \leq k$ for $f = \mu_i, \lambda_i, \kappa_p$ and each $k \geq 1$, is characterized by a finite list of forbidden minors.

5. Complete graphs and complete-graph bounds

For complete graphs we get better bounds than the general ones, but no definitive solution except in very small orders. Upper bounds on complete graphs imply bounds on simple graphs and more generally on multigraphs with bounded multiplicity.

First we discuss arbitrary representations.

Proposition 5.1. *$G(K_n, \varnothing)$ has no representation over $\mathbb{F}_q$ when $q \leq \lfloor \frac{n^2}{4} \rfloor$ if $n \geq 6$, when $q \leq 5$ if $n = 5$, when $q = 2$ if $n = 4$. If $n \geq 5$, $L(K_n, \varnothing)$ has no representation over $\mathbb{F}_q$ when $q < \lfloor \frac{n^2}{4} \rfloor$. $L(K_4, \varnothing)$ has no binary representation and $L(K_4^\circ, \varnothing)$ has no binary or ternary representation.*

Proof. From Lemma 2.4. □

Now we look into modular antivoltages.

Proposition 5.2. *We have $\mu_0(K_n) \geq \lfloor \frac{n^2}{4} \rfloor$ if $n \geq 3$.*

Proof. The bound follows from Proposition 3.1(a). □

In fact, it is easy to verify directly that $\mu_0(K_3) = 2$ and $\mu_0(K_4) = 4$; or see the discussion of μ_2 later.

Theorem 5.3. *For any simple graph Γ of order n, $\mu_2(\Gamma) \leq \lceil (e-2)(n-1)! \rceil$.*

Proof. We produce integral antivoltages for K_n whose breadth equals $1 + \lfloor (e-2)(n-1)! \rfloor$. Label the nodes $v_{-1}, v_0, v_1, \ldots, v_{n-2}$. For an edge $v_i v_j$ with $i < j$ the antivoltage is $\varphi_{ij} = (j-i)\pi_i$, where π_i equals the falling factorial $(n-1)_i = (n-1)(n-2)\cdots(n-i)$, with $(n-1)_0 = 1$, except that $\pi_{-1} = 0$. (Thus, the antivoltage of any edge at v_{-1} equals 0.) The important properties are that $\pi_{i-1} \mid \pi_i$ for $i > 0$ and that $\pi_i > \varphi_{jk}$ whenever $-1 \leq j < i$.

To prove no circle has antivoltage 0, let C be any circle and let v_i be the lowest-numbered node in C other than v_{-1}. Then $\varphi(C) \equiv (k-j)\pi_i \not\equiv 0 \pmod{\pi_{i+1}}$ if $C = v_j v_i v_k \cdots$ where $j, k \geq 0$, since $k - j < n - i$. A similar argument applies if v_{-1} is adjacent to v_i in C.

The Hamiltonian circle $H_0 = (v_{-1} v_0 v_1 \cdots v_{n-2})$ has antivoltage

$$1 + (n-1) + (n-1)_2 + \cdots + (n-1)_{n-3},$$

which by Taylor's remainder formula for e^x at $x = 1$ is equal to $\lfloor (e-2)(n-1)! \rfloor$. We show that H_0 is the unique circle with maximum antivoltage by building and demolishing bridges and levelling cliffs.

Take an arbitrary circle $C \neq H_0$, written as a cyclic permutation $(v_{i_0} v_{i_1} v_{i_2} \cdots v_{i_m})$ and oriented so that $\varphi(C) > 0$. We show how to modify C so as to increase $\varphi(C)$.

Case 1. $C = (v_i v_{-1} v_j \cdots)$ with $i < j$. By deleting v_{-1} from C we increase $\varphi(C)$ by $\varphi_{ij} > 0$.

Case 2. $v_{-1} \notin C$. There must be a decreasing step in C, that is, $C = (v_i v_j \cdots)$ where $i > j$. Insert v_{-1} between v_i and v_j. This increases $\varphi(C)$ by $\varphi_{ji} > 0$.

Case 3. $C = (v_{-1} v_{i_1} v_{i_2} \cdots v_{i_m})$ with $i_1 < i_m$. The structure of C is that it contains *peaks*, or locally maximal indices, and *valleys*, or locally minimal indices, joined by downslopes (from a peak to a valley) and upslopes. If C has more than one peak, it necessarily contains a *rising valley*, a valley whose preceding peak is lower than the succeeding peak.

Suppose C contains an ascending consecutive triple $v_i v_j v_k$ (that is, $-1 \leq i < j < k$). Then

$$\varphi_{ij} + \varphi_{jk} > \varphi_{ik}; \tag{5.1}$$

this follows because $(k-j)\pi_j > (k-j)\pi_i$. A descending consecutive triple satisfies the same inequality in reverse. Therefore any descending triple $v_k v_j v_i$ should be "bridged": v_j should be deleted and this will increase $\varphi(C)$. Thus we may assume henceforth that every downslope is a *cliff*, a peak followed immediately by a valley, as otherwise the antivoltage of C can certainly be increased.

Suppose C contains a peak v_p in a consecutive triple $v_i v_p v_j$ where $-1 \leq i < j$. Then $\varphi_{ip} + \varphi_{pj} = \varphi_{ip} - (p-j)\pi_j < 0$ since $\pi_j > \varphi_{ip}$. Because $\varphi_{ij} \geq 0$, we increase $\varphi(C)$ by deleting v_p (thus "levelling" the cliff $v_p v_j$). Hence we may assume that, if $C = (v_i v_p v_j \cdots)$ where v_p is a peak, then $i > j$.

Now consider a rising valley v_l. It must appear in a sequence $(v_i v_p v_l \cdots v_j v_k \cdots v_q \cdots)$ where v_p and v_q are the previous and next peaks (so $p < q$), $i > l \geq -1$, and $l \leq j < p < k \leq q$. We level the preceding peak by moving v_p to the upslope between v_j and v_k; i.e., C becomes $(v_i v_l \cdots v_j v_p v_k \cdots v_q \cdots)$. We show that this operation increases $\varphi(C)$ by comparing the parts of the sum that change: in C,

$$s = \varphi_{ip} + \varphi_{pl} + \varphi_{jk} = (p-i)\pi_i - (p-l)\pi_l + (k-j)\pi_j,$$

and in the modified circle,

$$s' = \varphi_{il} + \varphi_{jp} + \varphi_{pk} = -(i-l)\pi_l + (p-j)\pi_j + (k-p)\pi_p.$$

The difference is

$$s' - s = (k-p)(\pi_p - \pi_j) - (p-i)(\pi_i - \pi_l).$$

Since $i < p$, it follows that

$$s' - s \geq (\pi_p - \pi_j) - (p-i)\pi_i.$$

When $i > j$, $\pi_i > \pi_j$, so

$$\begin{aligned} s' - s &> \pi_p - (p+1-i)\pi_i \\ &= \pi_i[(n-1-i)_{p-i} - (p+1-i)] \\ &\geq \pi_i[(n-1-i) - (p+1-i)] \\ &= \pi_i[n-2-p] \geq 0. \end{aligned}$$

When $j > i$, $\pi_j > (p-i)\pi_i$, so we have

$$\begin{aligned} s' - s &> \pi_p - 2\pi_j \\ &\geq \pi_j[(n-j-1) - 2] = \pi_j[n-j-3] > 0, \end{aligned}$$

since $j \leq q-3 < n-3$. In every case, $s' > s$. Thus, moving v_p increases $\varphi(C)$.

What we have shown is that any rising valley can be eliminated while increasing the antivoltage. Consequently, we may assume C has only one peak. That means it has $i_1 < i_2 < \cdots < i_m$. Equation (5.1) shows that, if there is any place where $i_{k-1} < i_k - 1$, we increase the antivoltage by inserting $v_{i_{k-1}+1}$. Similarly, if $i_m < n-2$, we increase the antivoltage by inserting v_{n-2} at the end of the cycle. Thus in the end we have transformed C into H_0, which is therefore the only circle having maximum antivoltage. □

The upper bound for $n = 3$ gives $\mu_2 \leq \lceil 2(e-2) \rceil = 2$, which equals the lower bound. Hence,

$$\mu_0(K_3) = \mu_1(K_3) = \mu_2(K_3) = 2.$$

For $n = 4$ the upper bound is $\mu_2 \leq \lceil 6(e-2) \rceil = 5$. For $n = 5$ the bounds are $6 \leq \mu_0 \leq \mu_2 \leq \lceil 24(e-2) \rceil = 18$. In the upper bound, since $e - 2 \approx .71828$, we are getting about a 28% improvement over Theorem 3.4(a) by more cleverly assigning the antivoltages. The bound unfortunately appears to be far from best possible. K_4 regarded as a wheel (see Example 6.4) has integral antivoltages with $\max_C \varphi(C) = 3$; thus

$$\mu_0(K_4) = \mu_1(K_4) = \mu_2(K_4) = 4.$$

Moreover, K_5 has integral antivoltages shown in Figure 1 with $\max \varphi(C) = 8$. I found by a tedious calculation, which I believe is reliable, that there are no integral antivoltages whose breadth is 6 or 7. (This is the only example where I know $\mu_2 > \xi + 1$.) Thus,

$$8 \leq \mu_0(K_5) \leq \mu_2(K_5) \leq 9.$$

This is enormously better than what Theorem 5.3 implies. Indeed, I do not know whether μ_2 is linearly bounded either above or below by ξ for complete (or any) graphs, so the problem of effectually bounding $\mu_2(K_n)$ must be regarded as very open.

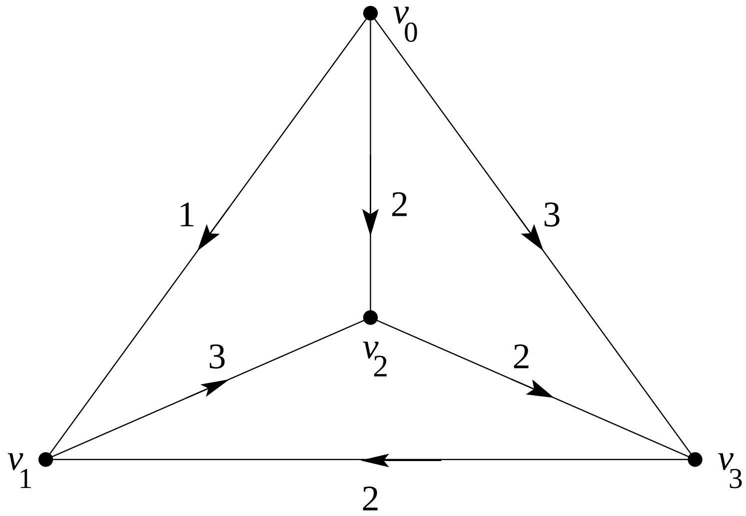

FIGURE 1. Integral antivoltages on K_5 in which the maximum antivoltage of a circle is $8 = \varphi(C)$ where $C = (v_{-1}v_0v_3v_1v_2)$. In the figure, v_{-1} and its 0-antivoltage edges are omitted.

Corollary 5.4. *If Γ is a loopless graph of order n and all edge multiplicities are at most m, then $\mu_2(\Gamma) \leq \lceil m^n(e^{1/m} - 1 - m^{-1})(n-1)! \rceil$.*

Proof. We exhibit an antivoltage for mK_n, which is K_n with each edge replaced by m copies of itself. In the proof of Theorem 5.3, replace each edge v_iv_j (where $0 \leq i < j$), whose antivoltage is $(j-i)\pi_i$, by m edges with antivoltages $\big[(j-i)m-k\big]m^i\pi_i$ for $k = 0, 1, \ldots, m-1$. The largest antivoltage of such an edge is $m^{i+1}(j-i)\pi_i$. The proof goes through if we assume all forward edges of C have this largest antivoltage. Thus $\varphi(C) \leq \varphi(H_0) = \sum_{i=0}^{n-3} m^{i+1}(n-1)_i$. □

There remain prime-power antivoltages; that is, upper and lower bounds on $\kappa_p(K_n)$ and $\lambda_i(K_n)$.

Proposition 5.5. *We have* $\kappa_p(K_3) = 1$,

$$\kappa_p(K_4) = 3,\ 2,\ 1 \text{ when } p = 2,\ p = 3,\ p \geq 5, \text{ and}$$

$$\kappa_p(K_n) \geq \log_p \lfloor n^2/4 \rfloor \text{ if } n \geq 5.$$

For $n \geq 3$ we have

$$\kappa_p(K_n) \leq \sum_{j=2}^{n-1} \lceil \log_p j \rceil \leq n - 2 + \log_p(n-1)!$$

but $\kappa_p(K_n) = 1$ if $p \geq (n-1)!$.

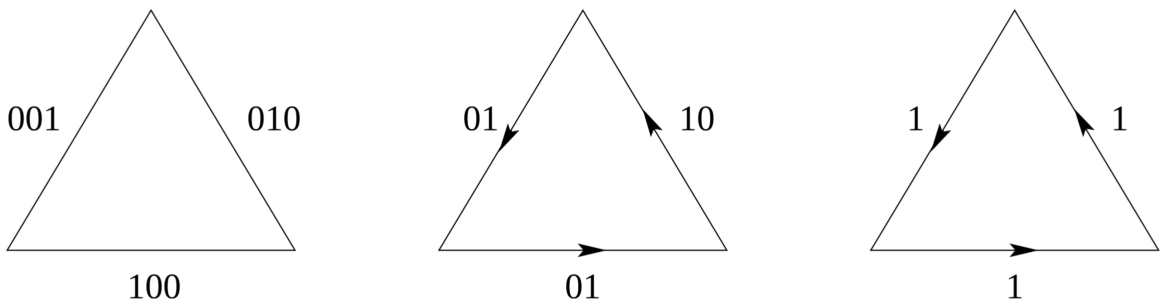

FIGURE 2. Antivoltages on K_4 with values in $\mathbb{Z}_p^k$ for (from left to right) $p^k = 2^3,\ 3^2$, and p^1 if $p \geq 5$. In the figures, one node and its 0-antivoltage edges are omitted.

Proof. Proposition 3.1(b) gives the lower bounds for $n \geq 5$. For the upper bound, in Theorem 3.4(b,b′) take T to consist of all edges at a node, so $\Gamma \setminus T = K_{n-1}$, $\sigma_p = \sum_{j=1}^{n-1} \lceil \log_p j \rceil$, and $\pi(\Gamma \setminus T, \alpha) = (n-1)!$. As for the exact values, the upper bounds are by example (trivial for K_3, and for K_4 see Figure 2). The lower bounds are based on the fact that by switching we can choose antivoltages to be zero on a node star; then the remaining edge (in K_3) or triangle (in K_4) must have antivoltages for which no path or circle sums to 0. □

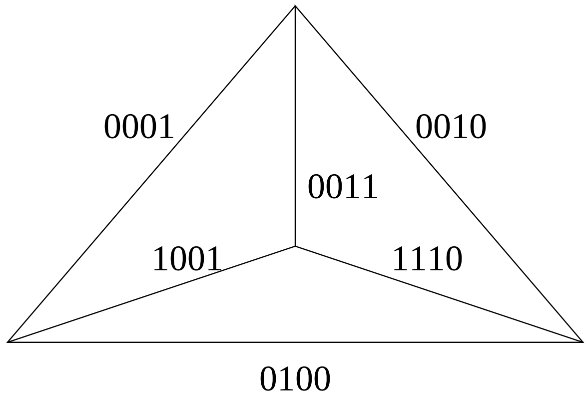

FIGURE 3. Antivoltages on K_5 with values in $\mathbb{Z}_2^4$. One node and its 0-antivoltage edges are omitted.

The value of $\kappa_2(K_5)$ is 4, strictly between the lower bound 3 and the upper bound 5. Figure 3 shows antivoltages in $\mathbb{Z}_2^4$. The proof that none exist in $\mathbb{Z}_2^3$ is neither difficult nor interesting, so we omit it.

Proposition 5.6. *We have*

$$\lambda_0(K_3) = \lambda_1(K_3) = 2, \quad \lambda_0(K_4) = \lambda_1(K_4) = 5,$$
$$6 \leq \lambda_0(K_5) \leq \lambda_1(K_5) < 4!^2 = 576,$$

and for $n \geq 6$,

$$\lfloor n^2/4 \rfloor \leq \lambda_0(K_n) \leq \lambda_1(K_n) < (n-1)!^2.$$

Proof. K_3 is obvious. K_4 follows from the analysis of κ_p. The cases $n \geq 5$ are based on Theorem 3.4 and Proposition 3.1. (Theorem 3.2 is much weaker.) □

Undoubtedly the upper bounds for $n \geq 5$ are not close. We set as a homework exercise to get a reasonable bound for K_5 by careful application of the degree method.

6. More examples

Here we treat a few more families of examples. In three we get a complete solution to the question of canonical bias representability although not necessarily lift representability.

Example 6.1 (*Frozen series-parallel graphs*). The problem of determining which fields admit representations of the bicircular or bicircular lift matroid is easy for certain series-parallel graphs that are, at bottom, fancy versions of Example 3.1. We define a *series-parallel graph* Γ as a loop, or a graph obtained from a link r_1r_2 by repetitions of the series and parallel operations of subdividing an edge and doubling an edge in parallel. The nodes r_1, r_2 are the *roots*; we regard a loop as the case in which $r_1 = r_2$. For an edge f in Γ, let P_f be an r_1r_2-path (or circle, in the loop case) that contains f. To get our restricted series-parallel graphs we confine the series and parallel operations to unfrozen edges, an edge f becoming *frozen* whenever an edge in $P_f \setminus f$ is doubled. (The initial link or loop r_1r_2 is unfrozen. The path P_f is unique for unfrozen edges in restricted graphs.) These restricted graphs (including the graph of a single loop), which we call *frozen series-parallel graphs*, are essentially the abstract duals of outerplanar graphs. (To be exact, frozen series-parallel blocks are the abstract duals of outerplanar blocks.[2] One can prove this by observing that the property of being a frozen series-parallel graph is preserved by Whitney 2-isomorphism operations. By such operations one can move all multiple edges to be incident to one of the roots. Then, in a planar drawing, this root is incident to every region; dually, the root becomes a region incident to every node. Thus we have an outerplanar block graph. This process is reversible. Since the forbidden minors for outerplanar graphs are K_4 and $K_{3,2}$ [2], by dualizing we find that a block is a frozen series-parallel graph if and only if it has no minor isomorphic to K_4 or $2C_3$.)

In a frozen series-parallel graph Γ, other than a loop, call two unfrozen edges *equivalent* if they lie in the same r_1r_2-path. Let m be the number of equivalence classes. We can produce antivoltages by assigning to one edge in each equivalence class, oriented from r_1 to r_2, a different antivoltage in the set $\{0, 1, \ldots, m-1\}$, and to all other edges antivoltage 0. This gives integral antivoltages with breadth m and hence antivoltages in each group $\mathbb{Z}_\mu$ with $\mu \geq m$. Furthermore, Γ contracts to mK_2 by contracting all the edges except one in each equivalence class. The number m of equivalence classes equals $\xi(\Gamma) + 1$. Therefore,

$$\mu_0(\Gamma) = \mu_1(\Gamma) = \mu_2(\Gamma) = \xi(\Gamma) + 1. \tag{6.1}$$

[2] I thank James Oxley for this observation.

Also, $\lambda_1(\Gamma)$ is the least prime power $q > \xi(\Gamma)$. These statements apply even when Γ is a loop. We draw the following conclusions:

Proposition 6.1. *Let Γ be a frozen series-parallel graph. $(\Gamma, \varnothing)$ has a canonical bias representation over $\mathbb{F}_q$ if and only if $q > \xi(\Gamma) + 1$. It has a canonical lift representation if and only if $q \geq \xi(\Gamma) + 1$. It has a lift or bias representation only if $q \geq \xi(\Gamma)$.*

Proof. The facts about canonical representations and the lower bounds on arbitrary representations follow as in Example 3.1 supplemented by Lemma 2.4, since $(mK_2, \varnothing)$ is a minor of $(\Gamma, \varnothing)$ and indeed $M(\Gamma) = M'(\Gamma) = M''(\Gamma) = m$. □

The existence of small noncanonical representations seems more difficult. Example 6.2 suggests that a bias representation may always exist if $q = \xi(\Gamma) + 1$, and that bias and lift representations sometimes exist if $q = \xi(\Gamma)$.

Example 6.2 (*Restricted multitriangles*). A frozen series-parallel block Γ of order 3 is K_3 with up to two multiple edges. (Then $L(\Gamma, \varnothing) = G(\Gamma, \varnothing)$.) We can determine the Galois fields $\mathbb{F}_q$ over which there is a representation of $L(\Gamma, \varnothing)$. Say there are edges e_{ij}^k, $1 \leq k \leq m_{ij}$, between v_i and v_j; thus v_1v_2 has multiplicity $m_{12} \geq 1$, v_2v_3 has multiplicity $m_{23} \geq 1$, and there is just one edge e_{13}.

Proposition 6.2. *When $m_{23} = 1$, a representation exists if and only if $q \geq \xi(\Gamma)$. When $m_{12}, m_{23} \geq 2$, a representation exists if and only if $q \geq \xi(\Gamma) + 1$.*

Proof. Comparing to Example 6.1, $m = m_{12} + m_{23} = \xi(\Gamma) + 1$. We have $q \geq m - 1$ from Proposition 6.1.

When $m_{23} = 1$, a representation is easy to construct for any $q \geq m - 1$.

Suppose $m_{12}, m_{23} \geq 2$ and assume a representation exists. Let $\hat{e}$ denote the point representing edge e. The line L_{12}, generated by the points $\hat{e}_{12}^i$ for $1 \leq i \leq m_{12}$, and the line L_{23}, defined similarly, meet in a point that represents no edge. The edge $\hat{e}_{13}$ lies on neither L_{12} nor L_{23}. On L_{23} there are: m_{12} points at which L_{23} intersects the lines $\hat{e}_{12}^i\hat{e}_{13}$, the point $L_{12} \wedge L_{23}$, and the m_{23} points $\hat{e}_{23}^j$. Therefore $q + 1 \geq m_{12} + m_{23} + 1$; so $q \geq m$. If $q = m$, the canonical lift representation shows that a representation exists. □

This example with $m_{12} > 1 = m_{23}$ is an instance where a representation of $G(\Gamma, \varnothing)$ or $L(\Gamma, \varnothing)$ is known to exist over a field that has no canonical bias or lift representation. When $m_{12}, m_{23} \geq 2$, every representation is obviously a canonical lift representation.

Example 6.3 (*Multitriangles*). Let Γ be a triangle with multiple edges, in which every edge is multiple. (Again, $L(\Gamma, \varnothing) = G(\Gamma, \varnothing)$.) By Theorem IV.7.1, $G(\Gamma, \varnothing)$ has only canonical lift and canonical bias representations. That reduces the question of representability over Galois fields entirely to the existence of modular and prime-power antivoltages. Because the bicircular and bicircular lift matroids are equal we deduce that, if Γ is a triangle with every edge multiple, then $G(\Gamma, \varnothing)$ is representable over $\mathbb{F}_q$ (where $q = p^k$) if and only if Γ has antivoltages in either $\mathbb{Z}_{q-1}$ or $\mathbb{Z}_p^k$.

Suppose the multiplicities are m_{12}, m_{23}, and m_{13}, their sum is m, and $m' = m - \min m_{ij}$. By Proposition 3.1, $\mu_0 \geq m'$ and $\lambda_0 \geq m'$. We show that $\mu_2 \leq m - 1 = \xi + 1$ by construction. With notation as in the preceding example, take integral antivoltages $\varphi(e_{12}^i) = i - 1$ for $1 \leq i \leq m_{12}$, $\varphi(e_{23}^j) = j - 1$ for $1 \leq j \leq m_{23}$, and $\varphi(e_{13}^k) = m_{12} + m_{23} - 2 + k$ for $1 \leq k \leq m_{13}$.

(We understand the subscripts to indicate the edge orientation.) Thus circles $e_{12}^i e_{23}^j e_{31}^k$ have antivoltages ranging between -1 and $-(m-2)$. We conclude that $\mu_2 \leq m-1 = \xi(\Gamma)+1$.

We leave κ_p and λ_1 as exercises.

Example 6.4 (*Wheels*). The t-spoke wheel W_t turns out to be simple, especially as regards the bicircular matroid. Note that t equals the cyclomatic number. We assume $t \geq 3$. Let the rim nodes be $x_1, x_2, \ldots, x_t$ in consecutive order and let the hub be z. Then

$$\mu_0(W_t) = \mu_1(W_t) = \mu_2(W_t) = t+1. \tag{6.2}$$

To prove $\mu_0 \geq t+1$, contract the path $x_1 x_2 \cdots x_{t-1}$ to a point, leaving a contrabalanced graph with three nodes and a $(t-1)$-fold multiple edge $x_1 z$. Now contract zx_t. This gives $((t+1)K_2, \varnothing)$. Apply Lemma 2.1 to deduce that $\mu_0 \geq t+1$. For $\mu_2 \leq t+1$, assign antivoltage 0 to each spoke and 1 to each rim edge $x_{i-1}x_i$ in that orientation (where $x_0 = x_t$). Clearly, these antivoltages are contrabalanced and the antivoltage of a circle reaches its maximum value t on the rim.

Proposition 6.3. *For $t \geq 3$, $G(W_t^\circ, \varnothing)$ has a representation over a field F if and only if $\#F \geq t+2$. No representation of $G(W_t, \varnothing)$ exists over any field of order $q < t$.*

For $t \geq 3$, $L(W_t^\circ, \varnothing)$ has a representation over a prime field $\mathbb{F}_p$ if and only if $p \geq t+1$. No representation of $L(W_t, \varnothing)$ exists over any field of order $q < t$.

Proof. Apply Equation (6.2) and Lemma 2.2. □

Antivoltages in $\mathbb{Z}_p^k$ are more difficult than modular antivoltages but Theorem 3.4 is helpful. We choose T so as to leave a large-degree node in $W_t \setminus T$ (I believe this is a good principle); that means T is a Hamiltonian path from the hub. Then $W_t \setminus T$ is $K_{1,t-1}$ with an extra edge hanging off one rim node, so $\pi(W_t \setminus T, \alpha) = 2t$ if the hub is the tail of every edge aside from the extra edge. Then the theorem gives

$$\mu_2 \leq 2t, \qquad \kappa_p \leq 1 + \lceil \log_p t \rceil, \qquad \lambda_1 < 4t^2. \tag{6.3}$$

When $t=3$, the true values are obtained by applying Propositions 5.5 and 5.6 since $W_3 = K_4$; we get $\kappa_p(W_3) = 3,2,1$ depending on p, and $\lambda_1(W_3) = 5$. The estimates of Theorems 3.2 and 3.4 are far too big; it seems likely that this remains true when p is large.

Example 6.5 (*Complete bipartite graphs*). For $K_{r,s}$ (with $r \geq s$), if we follow the advice of Example 6.4 about choosing T in Theorem 3.4 we will have $K_{r,s} \setminus T = K_{r-1,s-1}$. Orient all edges from one side to the other, so $d^o = (r-1, \ldots, r-1, 0, \ldots, 0)$ with $s-1$ positive terms. Then $\pi(K_{r,s} \setminus T, \alpha) = r^{s-1}$ and we get

$$\mu_2 \leq r^{s-1}, \qquad \kappa_p \leq (s-1)\lceil \log_p r \rceil, \qquad \lambda_1 < r^{2(s-1)}.$$

The opposite orientation gives $\pi = s^{r-1}$ instead, but Proposition 3.6 shows this is inferior.

7. The number of antivoltages

How many antivoltage mappings does Γ have with values in $\mathbb{Z}_\mu$? $\mathbb{Z}_p^k$? How many integral antivoltages have breadth bounded by a fixed positive integer (and are zero on a fixed maximal forest T; without this constraint the number is infinite)? For a finite abelian group $\mathfrak{A}$, let

$\alpha_\Gamma(\mathfrak{A}) =$ the number of antivoltages on Γ with values in $\mathfrak{A}$.

(The number that are zero on T equals $\alpha_\Gamma(\mathfrak{A})/(\#\mathfrak{A})^{\#T}$ because, with abelian antivoltages, each antivoltage that is zero on T switches uniquely to one with arbitrarily prescribed values on T.) Also, define

$\beta^\circ_{\Gamma,T}(\mu) =$ the number of integral antivoltages on Γ whose breadth is $\leq \mu$ and which are zero on T.

(The superscript circle on β is there for consistency with the notation in [1].) We cannot evaluate these functions very often but we can prove general properties. We need the cyclomatic number ξ of Γ, the rank $\rho = \#E - \xi$, and a convex polytope $P(\Gamma, T) \subseteq \mathbb{R}^{E\setminus T}$. Let x_C be the coefficient vector of the linear form $\varphi(C)$ corresponding to an oriented circle C; thus $\varphi(C) = \varphi \cdot x_C$. Think of $\mathbb{R}^{E\setminus T}$ as the subspace $\{x \in \mathbb{R}^E : x(e) = 0 \text{ for } e \in T\}$ of $\mathbb{R}^E$. Then $P(\Gamma, T)$ is the intersection of all the halfspaces

$$H_C = \{x \in \mathbb{R}^{E\setminus T} : x \cdot x_C \leq 1\}$$

corresponding to oriented circles C (so each circle gives two halfspaces). $P(\Gamma, T)$ is bounded, being contained in the box $[-1,1]^{E\setminus T}$; thus, its (ξ-dimensional) volume $\operatorname{vol} P(\Gamma, T)$ is finite, and the volume is positive since $P(\Gamma, T)$ is clearly full-dimensional in $\mathbb{R}^{E\setminus T}$.

Theorem 7.1. *$\alpha_\Gamma(\mathfrak{A})$ is a polynomial function of $\mu = \#\mathfrak{A}$, independent of the particular abelian group $\mathfrak{A}$. Its leading term is $\mu^{\#E}$ and it has a factor μ^ρ.*

The function $\beta^\circ_{\Gamma,T}(\mu)$ is a polynomial, independent of T, with leading term $\operatorname{vol} P(\Gamma, T) \cdot \mu^\xi$ *and nonzero constant term.*

Proof. Integral antivoltages first. Because each integral antivoltage switches to one of the same breadth that is 0 on T, $\beta^\circ_{\Gamma,T}(\mu)$ is independent of the choice of T. A vector $x \in \mathbb{R}^{E\setminus T}$ is an integral antivoltage with breadth $\leq \mu$ if and only if

$$\mu^{-1}x \in \Big(\operatorname{int} P(\Gamma, T) \setminus \bigcup \mathcal{H}\Big) \cap \mu^{-1}\mathbb{Z}^{E\setminus T},$$

where $\mathcal{H}$ is the set of all hyperplanes $x \cdot x_C = 0$ for $C \in \mathcal{C}$, the set of circles of Γ. In what we call in [1] "inside-out Ehrhart theory" (which is a combination of standard Ehrhart theory and poset Möbius inversion), $\beta^\circ_{\Gamma,T}(\mu)$ is the open Ehrhart quasipolynomial of $(P(\Gamma, T), \mathcal{H})$; thus it is a quasipolynomial in μ whose leading term is $\operatorname{vol} P(\Gamma, T) \cdot \mu^\xi$ and whose constant term is $\pm r$ where r is the number of regions into which $\mathcal{H}$ dissects $P(\Gamma, T)$. Moreover, $\beta^\circ_{\Gamma,T}(\mu)$ is a polynomial because all vertices of the regions are integral. To see this, note that any such vertex is determined by a subset of the linear equations in the matrix equation

$$\begin{bmatrix} I_T \ \ O \\ M(\mathcal{C}) \\ M(\mathcal{C}) \\ M(\mathcal{C}) \end{bmatrix} x = \begin{bmatrix} \mathbf{0} \\ \mathbf{1} \\ \mathbf{0} \\ \mathbf{-1} \end{bmatrix}, \tag{7.1}$$

where on the left the columns are indexed by E, I_T is an identity matrix, O is a zero matrix, and the rows of $M(\mathcal{C})$ are the vectors x_C^{T} for $C \in \mathcal{C}$ (without being duplicated for opposite orientations). $M(\mathcal{C})$ is well known to be totally unimodular, so any point determined by a subset of the rows of (7.1) is integral.

Similar reasoning applies to $\alpha_\Gamma(\mathfrak{A})$. For any subset $\mathcal{D}$ of $\mathcal{C}$, let $M(\mathcal{D})$ be the matrix that consists of the rows of $M(\mathcal{C})$ which are associated with $C \in \mathcal{D}$. Because $M(\mathcal{D})$ is totally unimodular, the number $\nu(\mathcal{D})$ of solutions of $M(\mathcal{D})x = \mathbf{0}$ equals $(\#\mathfrak{A})^{\#E - \operatorname{rk} M(\mathcal{D})}$. Thus, for instance, $\nu(\varnothing) = (\#\mathfrak{A})^{\#E}$. By inclusion and exclusion,

$$\alpha_\Gamma(\mathfrak{A}) = \sum_{\mathcal{D} \subseteq \mathcal{C}} (-1)^{\#\mathcal{D}} \nu(\mathcal{D}) = \sum_{\mathcal{D} \subseteq \mathcal{C}} (-1)^{\#\mathcal{D}} \mu^{\#E - \operatorname{rk} M(\mathcal{D})}.$$

There is a factor μ^ρ because $\#E - \operatorname{rk} M(\mathcal{D}) \geq \#E - \operatorname{rk} M(\mathcal{C}) = \#E - \xi = \rho$. □

We write $\alpha_\Gamma(\mu)$ instead of $\alpha_\Gamma(\mathfrak{A})$ for an abelian group with $\#\mathfrak{A} = \mu$, and $\beta^\circ_\Gamma(\mu) = \beta^\circ_{\Gamma,T}(\mu)$ for every T.

Note the corollary that $\operatorname{vol} P(\Gamma, T)$ is independent of T.

Now let us do some small examples.

Example 7.1. Let Γ consist of one node and m loops. Then

$$\alpha_\Gamma(\mu) = (\mu - 1)^m \quad \text{and} \quad \beta^\circ_\Gamma(\mu) = 2^m(\mu - 1)^m.$$

The same formulas hold if Γ consists of any m circles joined at one common node, except that α_Γ has to be multiplied by μ^ρ. For instance, if C is a circle of length l, then $\alpha_C(\mu) = \mu^{l-1}(\mu - 1)$ and $\beta^\circ_C(\mu) = 2(\mu - 1)$.

Example 7.2. The graph mK_2 consists of m links $e_1, \ldots, e_m$ joining two nodes. It is clear that $\alpha_{mK_2}(\mu) = (\mu)_m$. It is less obvious that

$$\beta^\circ_{mK_2}(\mu) = m(\mu - 1)_{m-1}. \tag{7.2}$$

For the proof, orient all edges the same way. Define $M =$ the set of μ-modular antivoltages, $B =$ the set of integral antivoltages with breadth $\leq \mu$,

$$M_i = \{\bar\varphi \in M : \bar\varphi(e_i) = 0\}, \qquad B_i = \{\varphi \in B : \varphi(e_i) = 0\},$$
$$B^+ = \{\varphi \in B : \varphi \geq 0 \text{ and } \min_j \varphi(e_j) = 0\}, \qquad B_i^+ = B_i \cap B^+.$$

Obviously, $\#B_1^+ = \cdots = \#B_m^+$ and $B^+ = B_1^+ \cup\!\!\!\cdot \cdots \cup\!\!\!\cdot B_m^+$, so $\#B^+ = m\#B_i^+$. The translation mappings $B \to B_i$ by $\varphi \mapsto \varphi - \varphi(e_i)$ and $B \to B^+$ by $\varphi \mapsto \varphi - \min_j \varphi(e_j)$ give bijections $B_i \leftrightarrows B^+$, so $\#B_i = \#B^+$. Modular translation $\bar\varphi \mapsto \bar\varphi - \bar\varphi(e_i)$ gives a μ-to-1 mapping $M \to M_i$, so $\#M_i = \mu^{-1}\#M$. Finally, suppose $\bar\varphi \in M_i$. We can interpret $\bar\varphi$ as a well-defined element of B_i^+ by treating $\mathbb{Z}_\mu = \{0, 1, \ldots, \mu - 1\}$ as a set of integers. This mapping is inverse to the modularization mapping $B_i^+ \to M_i$. Therefore, $\#M_i = \#B_i^+$. It follows that $\beta^\circ_{mK_2}(\mu) = \#B_i = \frac{m}{\mu}\#M = \frac{m}{\mu}\alpha_{mK_2}(\mu)$, thereby proving (7.2). We also see that the strong part of Problem 3.8 has a positive answer for the graphs mK_2.

The geometry of this example is that of the root system A_{m-1} in the projected form $A'_{m-1} = \{\pm b_i\} \cup \{b_j - b_i : j \neq i\} \subseteq \mathbb{R}^{m-1}$, where $\{b_i\}$ is the standard basis. If we assign $\varphi(e_i) = x_i$ for $i = 1, \ldots, m-1$, then $P(\Gamma, \{e_m\})$ is the polytope bounded by $x \cdot v \leq 1$ for all $v \in A'_{m-1}$; this is the polar dual polytope of $\operatorname{conv}(A'_{m-1})$. By (7.2), its volume is m. Taking $m = 3$ as an example, $\beta^\circ_{3K_2}(\mu)$ counts $\frac{1}{\mu}$-integral points in the plane domain bounded by $-1 < x_1 < 1$, $-1 < x_2 < 1$, $-1 < x_2 - x_1 < 1$ and not in the lines $x_1 = 0$, $x_2 = 0$, and $x_2 - x_1 = 0$.

I do not discern a deletion-contraction identity for either α_Γ or $\mu^{-\rho}\alpha_\Gamma$ or β°_Γ. Indeed, suppose any of them satisfied a recurrence of the form

$$F_\Gamma(\mu) = a(\mu)F_{\Gamma \setminus e}(\mu) + b(\mu)F_{\Gamma/e}(\mu)$$

when e is neither a loop nor an isthmus. Substituting $\Gamma = mK_2$ with $m = 2$ and $m = 3$ leads to inconsistent values for $b(\mu)$.

Example 7.3. I computed

$$\alpha_{K_4}(\mu) = \mu^3(\mu - 1)(\mu^2 - 3\mu + 3).$$

Contracting and deleting one edge,

$$\alpha_{K_4/e}(\mu) = \mu^2(\mu - 1)(\mu - 2)(\mu - 3), \qquad \alpha_{K_4\setminus e}(\mu) = \mu^2(\mu - 1)^2.$$

8. ROOT-OF-UNITY REPRESENTATIONS

Here is a final use for modular antivoltages. Call a *root-of-unity representation* of $(\Gamma, \varnothing)$ a canonical bias representation in which the antivoltages are in $\mathbb{Z}_\mu$ but we treat $\mathbb{Z}_\mu$ as the multiplicative group of powers of ζ, a complex μth root of unity; that is, we convert $\varphi(e) \in \mathbb{Z}_\mu$ to $\zeta^{\varphi(e)}$ before constructing the canonical bias representation. Thus we obtain a complex representation of $G(\Gamma, \varnothing)$ from a modular antivoltage mapping. Dualizing, we have a representation by an arrangement $\mathcal{H}[\Gamma, \varphi]$ of complex hyperplanes. All the machinery of complex arrangements can be applied [6]; in particular, the Poincaré polynomial of the complement is determined by Γ since it is a simple transform of the chromatic polynomial of $(\Gamma, \varnothing)$.

REFERENCES

[1] Matthias Beck and Thomas Zaslavsky, Inside-out polytopes, *Advances in Math.* **205** (1) (2006), 134–162. MR 2007e:52017.

[2] Gary Chartrand and Frank Harary, Planar permutation graphs, *Ann. Inst. H. Poincaré Sect. B (N.S.)* **3** (1967), 433–438. MR 37 #2626.

[3] S.M. Gagola, Solution to Problem 10606, Avoiding the identity, *Amer. Math. Monthly* **106** (6) (June–July, 1999), 590–591.

[4] Victor Klee, The greedy algorithm for finitary and cofinitary matroids, in: Theodore S. Motzkin, ed., *Combinatorics*, pp. 137–152, Proc. Symp. Pure Math., Vol. 19, Amer. Math. Soc., Providence, R.I., 1971. MR 48 #10865.

[5] L.R. Matthews, Bicircular matroids, *Quart. J. Math. Oxford* (2) **28** (1977), 213–227. MR 58 #21732.

[6] Peter Orlik and Hiroaki Terao, *Arrangements of hyperplanes*, Grundl. math. Wiss., Vol. 300. Springer-Verlag, Berlin, 1992. MR 94e:52014.

[7] J.G. Oxley, *Matroid Theory*, Oxford Univ. Press, Oxford, 1992. MR 94d:05033.

[8] M.J. Piff and D.J.A. Welsh, On the vector representation of matroids, *J. London Math. Soc.* (2) **2** (1970), 284–288. MR 41 #3308.

[9] N. Robertson and P.D. Seymour, Graph minors—a survey, in: I. Anderson, ed., *Surveys in Combinatorics 1985* (Proc. Tenth British Combin. Conf., Glasgow, 1985), pp. 153–171, London Math. Soc. Lecture Note Ser., 103, Cambridge Univ. Press, Cambridge, 1985. MR 87e:05130.

[10] ——, Graph minors. XX. Wagner's conjecture, *J. Combin. Theory Ser. B* **92** (2) (2004), 325–357. MR 2005m:05204.

[11] J.M.S. Simões-Pereira, On subgraphs as matroid cells, *Math. Z.* **127** (1972), 315–322. MR 47 #6522.

[12] ——, On matroids on edge sets of graphs with connected subgraphs as circuits II, *Discrete Math.* 12 (1975), 55–78. MR 54 #7298.

[13] Thomas Zaslavsky, Bicircular geometry and the lattice of forests of a graph, *Quart. J. Math. Oxford* (2) **33** (1982), 493–511. MR 84h:05050c.

[14] ——, Geometric lattices of structured partitions. I: Gain-graphic matroids and group valued partitions. Manuscript, 1985.

[15] ——, Biased graphs. I. Bias, balance, and gains. II. The three matroids. III. Chromatic and dichromatic invariants. IV. Geometrical realizations, *J. Combin. Theory Ser. B* **47** (1989), 32–52; **51** (1991), 46–72; **64** (1995), 17–88; **89** (2003), 231–297. MR 90k:05138; 91m:05056; 96g:05139; 2005b:05057.

[16] ——, Frame matroids and biased graphs. *European J. Combin.* 15 (1994), 303–307. MR 95a:05021.

Binghamton University, Binghamton, New York 13902-6000