Signed Analogs of Bipartite Graphs

Thomas Zaslavsky
Binghamton University
Binghamton, New York, U.S.A. 13902-6000

September 1991

Revised July 1996

Abstract

We characterize the edge-signed graphs in which every "significant" positive closed walk (or combination of walks) has even length, under seven different criteria for significance, and also those edge-signed graphs whose double covering graph is bipartite. If the property of even length is generalized to positivity in a second edge signing, the characterizations generalize as well.

We also characterize the edge-signed graphs with the smallest nontrivial chromatic numbers.

Research and preparation of this report were supported by the SGPNR under grant 9009BCW747 and by the NSF under grant DMS-9200472.

Many thanks to the referee for helping me make this paper readable.

A problem in topological graph theory led me to investigate how the characterizations of bipartite graphs by even length of closed walks or even length of polygons (graphs of simple closed paths) extend to finite *signed graphs* (graphs with signed edges). A bipartite graph can be characterized as a graph whose face boundaries, in any surface embedding, all have even length. I wanted to know which signed graphs have this *even embedding property*. The embedding rule I used is that of *orientation embedding*, where the graph is embedded so that a polygon, or equivalently any closed walk, preserves surface orientation precisely when the product of its edge signs is positive. Because a face boundary, regarded as a closed walk, is positive, and every positive polygon is a face boundary in some embedding, the even embedding property of a signed graph is intermediate between the properties of evenness of positive closed walks and of positive polygons. Thus it was natural to ask which signed graphs have either of the latter properties.

In answering these questions the last property, the evenness of positive polygons, turns out to be fundamental, for it already implies the other two properties for loopless block graphs. The properties do differ slightly, however, for other connected graphs, so things are more complicated than with bipartiteness of unsigned graphs.

Having gone thus far, one can generalize evenness of polygons in several other ways. For instance, regarding a polygon as a matroid circuit leads one to investigate even length of three kinds of subgraphs that are circuits in matroids associated with a signed graph. Orientation embedding suggests asking about bipartiteness of the signed covering graph (to be defined later on). And finally, to show that a graph is bipartite one need not verify even length of all polygons; chordless polygons suffice. We extend this observation by finding the appropriate signed-graphic analog of a chordless polygon. (It is not, as one might think, a positive chordless polygon. Our work does not let us determine the signed graphs in which every such polygon has even length. That is a much deeper question, but it ought to be answerable by means of Truemper's description in [**7**] of the classes of chordless polygons in a graph that can be the positive ones in some signing.)

Bipartite graphs are usually defined as, in essence, the graphs with the smallest nontrivial chromatic number—that is, as the bicolorable graphs. This definition too can be generalized (quite easily) to signed graphs but one does not obtain quite the same classes of graphs as with the evenness and covering properties we mentioned initially. For the results see the final section.

Definitions

We denote by $\Sigma = (\Gamma, \sigma)$ a finite signed graph, where $\Gamma = (V, E)$ is the underlying graph and $\sigma : E \to \{+, -\}$ the sign labelling. Loops and multiple edges are allowed. The sign of a walk W whose edge sequence is $e_1, e_2, \cdots, e_l$ is $\sigma(W) = \sigma(e_1)\sigma(e_2)\cdots\sigma(e_l)$. Σ is *bipartite* if Γ is bipartite. It is *balanced* if every polygon is positive and *antibalanced* if negating its signs makes it balanced—that is, if every even polygon is positive and every odd one is negative. A *block* is a connected graph with no cutpoint. We regard a vertex supporting a loop and incident with another edge as a cutpoint, so a block is either the graph of a loop or is a loopless 2-connected graph. A *block of* Σ is a maximal block subgraph.

A *chord* of a polygon C is an edge whose endpoints are distinct vertices of C that are nonadjacent in C. In a signed graph, a *balanced chord* of a positive polygon C is a chord e for which $C \cup e$ is balanced. A *pit* is a positive polygon that has no balanced chord.

Switching σ means reversing the sign of every edge that has one endpoint in a vertex subset X and the other in its complement. We call σ and σ' *switching equivalent* if σ' is obtained from σ by switching; this is obviously an equivalence relation. Clearly, Σ is balanced if and only if it switches to all positive (this is essentially Harary's characterization of balance [**1**]); if X is a suitable switching set we call $\{X, V \backslash X\}$ a *Harary bipartition* of V.

In the two principal matroids of Σ, the *bias* and *lift* matroids [**8**], among the circuits are the positive polygons and the unions of pairs of negative polygons whose intersection is a single vertex. These are called *tight circuits.* Each tight circuit is a face boundary in some embedding of Σ (by Lemma 7 below, part (a) and part (b) with $k = 0$). That is not the case in general for the remaining circuits in either matroid, which we call *loose.* A loose lift circuit is a pair of vertex-disjoint negative polygons; a loose bias circuit is the same together with a minimal connecting path. The *length* of a matroid circuit C is the number of edges in it, except that if C is a loose bias circuit, then its length is the number of edges in its two polygons plus twice the number in the connecting path. The *sign* $\sigma(C)$, if C is not a polygon, equals the product of the signs of the two polygons it contains. (These definitions are based on the fact that a shortest closed walk — or in the case of a loose lift circuit, a shortest pair of closed walks — that covers C will traverse each polygon once and a connecting path twice.)

The *signed covering graph* $\tilde{\Sigma}$ is an unsigned graph whose vertex set is $\tilde{V} = \{+, -\} \times V$ and whose edge set consists of two edges, let us say $\tilde{e}_+$ and $\tilde{e}_-$, for each $e \in E$, such that the endpoints of $\tilde{e}_\delta$ are (δ, v) and $(\delta\sigma(e), w)$ if those of e are v and w. (The subscript in $\tilde{e}_\delta$ is simply a notational device, not an intrinsic sign.) Setting $p(\delta, v) = v$ and $p(\tilde{e}_\delta) = e$ defines the *covering projection* $p : \tilde{\Sigma} \to \Sigma$. A *fiber* is a set $p^{-1}(x)$ where $x \in V \cup E$. Let α denote the unique fiber-preserving, fixed-point-free automorphism of $\tilde{\Sigma}$. A walk W in Σ *lifts* (not uniquely) to a walk $\tilde{W}$ in $\tilde{\Sigma}$ by the rule that $\tilde{W}$ is a walk and $p(\tilde{W}) = W$. Then $\tilde{W}$ is closed precisely when W is closed and positive. W lifts to precisely two walks; each is the image under α of the other.

It is often possible to regard the property of even length as merely the particular instance of signed-graph positivity in which the edge signs are all negative, and to generalize evenness results to signed graphs without substantial change in the proofs. That is the case here. Thus we introduce a second edge signing σ_2, writing $\Sigma_2 = (\Gamma, \sigma_2)$, and we generalize our results to *doubly signed graphs* (Σ, σ_2). We define $\tilde{\sigma}_2$ to be the signing of $E(\tilde{\Sigma})$ given by $\tilde{\sigma}_2(\tilde{e}) = \sigma_2(p(\tilde{e}))$. We call a walk *positive* if it is σ-positive. We call (Σ, σ_2) *switchable* if $(\Gamma, \sigma\sigma_2)$ is balanced, *doubly balanced* if Σ and Σ_2 are both balanced. We say Σ has the *σ_2-positive embedding property* if every face boundary walk in every embedding of Σ is σ_2-positive. If we take σ_2 to be all negative, then σ_2-positivity becomes even length, σ_2-balance becomes bipartiteness, and switchability becomes antibalance.

Evenness: The Theorems

We state first the general theorem, then as a corollary the special case where σ_2 is all negative, which contains as part (iii) the embedding theorem in which I was originally interested. Then we give the proofs.

Theorem 1. *Let (Σ, σ_2) be a doubly signed graph.*

(i) The following three statements are equivalent.
  - *(a) All pits (of Σ) are σ_2-positive.*
  - *(b) All positive polygons are σ_2-positive.*
  - *(c) Each block of Γ is σ_2-balanced or switchable.*

(ii) All tight circuits in Σ are σ_2-positive $\Leftrightarrow$ after deleting doubly balanced blocks, each connected component is σ_2-balanced or switchable.

(iii) Σ has the σ_2-positive embedding property $\Leftrightarrow$ after deleting all isthmi of Γ, each component is σ_2-balanced or switchable.

(iv) The following properties of (Σ, σ_2) are equivalent:
  - *(a) All bias circuits of Σ are σ_2-positive.*
  - *(b) All positive closed walks are σ_2-positive.*
  - *(c) $(\tilde{\Sigma}, \tilde{\sigma}_2)$ is balanced.*
  - *(d) Each connected component of Γ is σ_2-balanced or switchable.*

(v) All lift circuits in Σ are σ_2-positive $\Leftrightarrow$ (Σ, σ_2) is σ_2-balanced or switchable.

Theorem 2. *Let Σ be a signed graph.*

(i) The following three statements are equivalent.
  - *(a) Every pit has even length.*
  - *(b) Every positive polygon has even length.*
  - *(c) Each block is bipartite or antibalanced.*

(ii) All tight circuits have even length $\Leftrightarrow$ after deleting balanced bipartite blocks, each connected component is bipartite or antibalanced.

(iii) Σ has the even embedding property $\Leftrightarrow$ after deleting all isthmi of Γ, each component is bipartite or antibalanced.

(iv) The following properties of Σ are equivalent:
  - *(a) All bias circuits have even length.*
  - *(b) All positive closed walks have even length.*
  - *(c) $\tilde{\Sigma}$ is bipartite.*
  - *(d) Each connected component is bipartite or antibalanced.*

(v) All lift circuits have even length $\Leftrightarrow$ Σ is bipartite or antibalanced.

The import of Theorem 2 is that the signed graphs most closely analogous to bipartite unsigned graphs are the antibalanced and bipartite ones. It seems odd that the even embedding property of Σ is not equivalent to that of $\tilde{\Sigma}$ (i.e., bipartiteness of $\tilde{\Sigma}$, since $\tilde{\Sigma}$ is unsigned.)

Evenness: The Lemmas

Lemma 1. *A signed graph in which every chordless polygon is positive, is balanced.*

Proof. Consider a shortest negative polygon C, if one exists. It has a chord, hence is the symmetric difference of two shorter polygons, both necessarily positive. Therefore C is positive, contradicting the assumed negativity of C. ■

The statement and proof readily generalize to biased graphs (which are defined in [**8**]).

Lemma 2. *Given signings σ and σ' of Γ, properties (a)-(d) are equivalent.*

(a) σ and σ' are switching equivalent.

(b) $(\Gamma, \sigma\sigma')$ is balanced, where $\sigma\sigma'(e)$ means $\sigma(e)\sigma'(e)$.

(c) $\sigma(W) = \sigma'(W)$ for every closed walk W.

(d) $\sigma(C) = \sigma'(C)$ for every polygon C.

Proof. (a) $\Rightarrow$ (c) $\Rightarrow$ (d): Trivial.

(a) $\Leftrightarrow$ (b): There is a set X such that switching X converts σ' to σ $\Leftrightarrow$ there is an X whose switching converts $\sigma\sigma'$ to all positive $\Leftrightarrow$ $(\Gamma, \sigma\sigma')$ is balanced.

(b) $\Leftrightarrow$ (d): $(\Gamma, \sigma\sigma')$ is balanced $\Leftrightarrow$ $\sigma\sigma'(C) = +$ for every polygon C $\Leftrightarrow$ $\sigma(C) = \sigma'(C)$ for every C. ∎

Lemma 3. (Harary [**2**]) *If Σ is an unbalanced signed block, then every vertex lies on a negative polygon.*

Proof. (This proof differs from Harary's.) Let C be a negative polygon and v a vertex not in C. We apply the vertex Menger theorem to produce a negative polygon on v. Σ, being a block of order exceeding one, can have no loops, so C has order at least two. Consequently, if we add a new vertex x adjacent to every vertex of C, we still have a block. By Menger's theorem there are two internally disjoint paths from v to x. The portions of these paths from v to C, together with C itself, form a theta graph (a subdivision of a triple edge) in Σ. Call the three paths composing the theta graph P, Q, and R, labelled so that $P \cup Q = C$ and v lies in R. Both $P \cup R$ and $Q \cup R$ are polygons, and exactly one is negative, because $\sigma(P)\sigma(Q) = \sigma(C) = -$. Thus v lies on a negative polygon. ∎

The statement of Lemma 3 generalizes directly to biased graphs. The proof also generalizes with a slight adaptation.

Lemma 4. *If Γ is a block which contains a Σ_2-negative lift circuit, then it contains a Σ_2-negative but positive polygon.*

Proof. A σ_2-negative lift circuit that is not itself a positive polygon consists of two negative polygons, C_1 and C_2, such that (with suitable notation) $\sigma_2(C_1) = -$ and $\sigma_2(C_2) = +$. Much as in the proof of Lemma 3, we deduce, by adding vertices x_1 and x_2 adjacent respectively to every vertex of C_1 and C_2, that there are two vertex-disjoint paths R and S from C_1 to C_2, internally disjoint from C_1 and C_2. (If C_1 and C_2 share a vertex, we can and do take R to be the path of length zero at that vertex.) The endpoints of R and S on C_1 divide C_1 into paths P_1 and Q_1, which have opposite signs in Σ_2. Similarly, C_2 divides into P_2 and Q_2, but their signs in Σ_2 are equal. Of the polygons $D = P_1 \cup R \cup S \cup P_2$ and $Q_1 \cup R \cup S \cup P_2$, one is σ_2-negative, say the former. Thus $D' = P_1 \cup R \cup S \cup Q_2$ is also σ_2-negative, while its σ-sign is opposite to that of D. Therefore D or D' is a positive polygon with negative sign in Σ_2. ∎

Lemma 5. *Any positive but σ_2-negative closed walk W in Σ contains within its edges a σ_2-negative lift circuit.*

Proof. Let E' be the set of edges which are traversed by W an odd number of times. This set has all even degrees; hence it is the edge-disjoint union of polygons $D_1, D_2, \ldots, D_q$.

Choose such a decomposition into the most possible polygons. Now, $\sigma(W)=\prod\sigma(D_i)$ so an even number of the polygons D_i are negative. At the same time $\sigma_2(W)=\prod\sigma_2(D_i)$ so an odd number of D_i are σ_2-negative; we are done unless all of them are negative. In that case E' contains two negative polygons, a σ_2-negative one C_1 and a σ_2-positive one C_2, which are edge disjoint.

If C_1 and C_2 have at most one common vertex, then they form a σ_2-negative lift circuit. If they have exactly two vertices in common, they form a subdivision of a quadruple edge from which it is easy to construct a σ_2-negative positive polygon. To complete the proof we show that C_1 and C_2 cannot share more than two vertices. If there were three vertices of intersection, say x_1, x_2, and x_3, dividing C_1 into thirds P_{12}, P_{23}, P_{31} and C_2 into Q_{12}, Q_{23}, Q_{31}, then $P_{12}\cup Q_{12}$, etc., would be even-degree subgraphs, collectively decomposable into at least three edge-disjoint polygons. That is contrary to the assumption. So the lemma is proved. ■

Lemma 6. *Let (Σ,σ_2) be a connected doubly signed graph for which $(\tilde\Sigma,\tilde\sigma_2)$ is balanced, and let $\{X,\tilde V\backslash X\}$ be a Harary bipartition of $\tilde V$ in $(\tilde\Sigma,\tilde\sigma_2)$. Then either X is a union of vertex fibers or it consists of one vertex from each vertex fiber. If the former, (Σ,σ_2) is switchable. If the latter, it is σ_2-balanced.*

Proof. (This is the only proof that is simpler when $\sigma_2\equiv -$.) Let $\tilde W$ be a walk in $\tilde\Sigma$, say from (δ,v) to (ε,w), so that $\alpha(\tilde W)$ is a walk from $(-\delta,v)$ to $(-\varepsilon,w)$. Evidently, if $\tilde\sigma_2(\tilde W)=+$, then $(\delta,v)\in X \Leftrightarrow (\varepsilon,w)\in X$, while if $\tilde\sigma_2(\tilde W)=-$, then $(\delta,v)\in X\Leftrightarrow (\varepsilon,w)\notin X$. It is also clear that $\tilde\sigma_2(\tilde W)=\sigma_2(p(\tilde W))=\tilde\sigma_2(\alpha(\tilde W))$. From this it is easy to see that $X\cap p^{-1}(v)$ and $X\cap p^{-1}(w)$ have the same parity. Thus, either X is a union of vertex fibers, or $|X\cap p^{-1}(v)|=1$ for every v.

In the former case a closed path P at $v\in V$ lifts to a path $\tilde P$ from $(+,v)$ to itself or to $(-,v)$. Since $(\tilde\Sigma,\tilde\sigma_2)$ is balanced with $(+,v)$ and $(-,v)$ in the same half of the Harary bipartition, $\tilde P$ is positive, whence P is also positive. Therefore Σ_2 is balanced.

In the latter case we may reverse the signs in certain vertex fibers of $\tilde\Sigma$ so as to make $X=\{+\}\times V$; this amounts to switching the corresponding vertices of Σ. The new Σ equals Σ_2 so the original was switching equivalent to it. ■

Lemma 7. *(a) Let C be a positive polygon in a signed graph Σ. Then a walk once around C is a face boundary walk in some embedding of Σ. (b) Let $C_0, C_1, \dots, C_{k+1}$ be polygons in a signed graph Σ such that each $C_{i-1}\cap C_i$ is a vertex v_i, $C_j\cap C_i=\emptyset$ for $i>j+1$, C_0 and C_{k+1} are negative, and $C_1,\dots,C_k$ are positive. Let the two paths in C_i from v_i to v_{i+1} be P_i and Q_i. Then $W=C_0P_1\cdots P_kC_{k+1}Q_k^{-1}\dots Q_1^{-1}$ is a face boundary walk in some embedding of Σ, where Q_i^{-1} denotes Q_i taken in the reverse direction, and C_i (for $i=0,k+1$) here denotes a walk around the polygon C_i with starting point v_1 or v_{k+1}, as appropriate.*

Proof. We begin with definitions from [**9**, Section 6], slightly simplified. Because of the simplification, in the proof we need to subdivide loops so that no edge has both ends at the same vertex.

A *rotation* at a vertex is a cyclic permutation of the edges at the vertex. A *rotation system* $\mathbf{R}$ for a signed graph is an assignment of a rotation $\mathbf{R}(v)$ to each vertex. Switching a vertex v has the effect of replacing $\mathbf{R}(v)$ by its inverse.

Rotation systems for signed graphs are identical or equivalent to Ringel's combinatorial maps [**5**], Stahl's "generalized embedding schemes" [**6**], and Lins' "graph-encoded maps" [**3**, **4**].

A *local orientation* of a rotation system $\mathbf{R}$ for Σ is a choice at a particular vertex v of one of $\mathbf{R}(v)$ or $\mathbf{R}(v)^{-1}$, that is of $\mathbf{R}(v)^{\epsilon}$ where $\epsilon = +1$ or -1. (More simply, it is a choice of $\epsilon = +1$ or -1.) Given a local orientation ϵ_0 at v_0 (the *initial orientation*) and a walk $W = v_0 e_1 v_1 \ldots e_k v_k$, the *transported orientation* at v_i is $\mathbf{R}(v_i)^{\sigma(e_1 e_2 \ldots e_i)\epsilon_0}$, briefly $\epsilon_i = \sigma(e_1 e_2 \ldots e_i)\epsilon_0$. We call W together with an initial orientation and all transported orientations an *oriented walk*. Here is a rule we may use to take an oriented walk on Σ, given $\mathbf{R}$. Start at any vertex v_0 with local orientation ϵ_0 and take any edge e_1 at v_0. This leads to a vertex v_1. After i steps we will have taken an oriented walk $W_i = (v_0, e_1, v_1, \ldots, e_i, v_i)$ with initial orientation ϵ_0. The next edge in the walk is that which succeeds e_i in the permutation $\mathbf{R}(v_i)^{\sigma(W_i)\epsilon_0}$. The walk terminates when it returns to v_0 with positive sign $\sigma(W_i)$ and the next edge e_{i+1} would equal e_1. Such an oriented walk we call an $\mathbf{R}$-*walk*. Suppose Σ is embedded; we say an $\mathbf{R}$-walk W *bounds* a face F if F is always on the right in the local orientation as we take the walk. (That is, F is the face separating the incoming and outgoing edges at each vertex of W.) By [**5**, Theorem 12] or (essentially) [**6**, Theorem 2] (or consult [**9**, Lemmas 6.1 and 6.2]), given $\mathbf{R}$ there is an embedding of Σ in which every $\mathbf{R}$-walk bounds a face.

Now we construct a rotation system $\mathbf{R}$ for Σ that makes the walk W of the lemma an $\mathbf{R}$-walk. In (b) we need some notation: at v_i, let the edge ends in C_{i-1} be e'_i and f'_i and those in C_i be e_i and f_i, named so $e_i, e'_{i+1} \in P_i$ and $f_i, f'_{i+1} \in Q_i$.

Switch so all edges of W are positive except, in case (b), f'_1 and f_k. Choose $\mathbf{R}(v_i)$ to have the form $(e_i e'_i \ldots f'_i f_i \ldots)$. For each divalent vertex v in W, whose edges in W are e and f in the order traced on W, let $\mathbf{R}(v) = (ef \cdots)$ except that, in (b), $\mathbf{R}(v) = (fe \cdots)$ if v is in a Q_i. For all other vertices u in Σ, $\mathbf{R}(u)$ is arbitrary. Then W is an $\mathbf{R}$-walk, hence bounds a face in some embedding of Σ. ■

Lemma 8. *Let Σ be embedded in a surface. Suppose that a face boundary walk W has a repeated edge e and has the form $W = eW_1e^{\pm 1}W_2$ (where e^{-1} means e traversed in the opposite direction). Then W_1 and W_2 have the same sign, positive if $W = eW_1e^{-1}W_2$ and negative if $W = eW_1eW_2$.*

Proof. First of all, $\sigma(W_1) = \sigma(W_2)$ because a face boundary walk is positive and $\sigma(W) = \sigma(e)\sigma(W_1)\sigma(e)\sigma(W_2)$.

Let F be the face bounded by W. Suppose W is directed so that, as we trace it, F lies on the right. For e to be repeated in W, two different parts of F must be on the two sides of e and the two passages of e must have these two parts of F on the right. If W_1 is positive, we return to e with preserved orientation. To have a different part of F on the right, we must be traversing e in the direction opposite to the first traversal. If W_2 is negative, so that we return to e with reversed orientation, we must trace e the second time in the same direction as in the first traversal. ■

Lemma 9. *Let T be a tree with signed vertices in which the number of negative vertices is positive and even. There is an edge e such that each component of $T \backslash e$ has an odd number of negative vertices.*

Proof. We do induction on n, the order of T. For $n = 2$, the lemma is trivial. For larger order, consider $T \backslash f$ where f is any edge of T. Let T_1 and T_2 be the components of $T \backslash f$. If a T_i has oddly many negative vertices, we take $e = f$. Otherwise, say T_1 has some negative vertices. By induction T_1 has an edge e such that each component of $T_1 \backslash e$ has an odd number of negative vertices. One of these components is a component of $T \backslash e$. Thus the lemma is proved. ■

Evenness: The Proof

We begin at the beginning with part (i). The implications (c) $\Rightarrow$ (b) $\Rightarrow$ (a) are trivial.

Part (i), (a) $\Rightarrow$ (b): Assume every pit is σ_2-positive. We prove first that a maximal balanced subgraph Γ' is σ_2-balanced. Every chordless polygon in Γ' is a pit in Σ, hence σ_2-positive. Thus Γ' is σ_2-balanced, by Lemma 1.

Now look at any positive polygon C. Extend it to a maximal balanced subgraph. Since the latter is σ_2-balanced, so is C.

Part (i), (b) $\Rightarrow$ (c): We may suppose Γ is a block. For the moment identify $+$ and $-$ with 0 and 1 in $\mathbb{F}_2$. Let Z be the cycle space of Γ, i.e., the subspace of $\mathbb{F}_2^E$ spanned by the characteristic vectors of polygons. Both σ and σ_2 determine linear functionals on $\mathbb{F}_2^E$, hence on Z; call these functionals respectively ϕ and ϕ_2: $Z \to \mathbb{F}_2$. Since $\phi^{-1}(0)$ has codimension at most 1, and since the hypothesis implies that $\phi^{-1}(0) \subseteq \phi_2^{-1}(0)$, $\phi_2^{-1}(0)$ can only be $\phi^{-1}(0)$ or Z. In the former case (Σ, σ_2) is switchable. In the latter it is σ_2-balanced.

Part (ii), $(\Rightarrow)$: We may suppose (Σ, σ_2) without doubly balanced blocks. Assume all tight circuits are σ_2-positive. From (i) we see that each block is σ_2-balanced or switchable. If there is a component that is neither σ_2-balanced nor switchable, then there are an unbalanced σ_2-balanced block Γ_1 and an unbalanced switchable block Γ_2 that have a common vertex v. By Lemma 3 applied to Σ, there are negative polygons C_1 on v in Γ_1 and C_2 on v in Γ_2. Then $C_1 \cup C_2$ is a σ_2-negative tight circuit, contrary to assumption.

Part (ii), $(\Leftarrow)$: Again we may suppose (Σ, σ_2) without doubly balanced blocks. Assume each component is σ_2-balanced or switchable, so there are no σ_2-negative positive polygons. Consider a tight circuit C. If C is contained in a block, Lemma 4 forces it to be σ_2-positive. If C is not contained in a block, then $C = C_1 \cup C_2$ where C_1 and C_2 are negative polygons that lie in different blocks and meet at a vertex. The blocks are in the same component, hence are both switchable or both σ_2-balanced. Either way, C is σ_2-positive.

Part (iii). We let $\Gamma' = \Gamma \backslash \{\text{isthmi}\}$ and $\Gamma'_1, \Gamma'_2, \cdots$ be the components of Γ'.

$(\Rightarrow)$: Assume Σ has the σ_2-positive embedding property. Because of Lemma 7(a), every positive polygon is σ_2-positive. Thus by part (i) every block is switchable or σ_2-balanced.

Suppose a component of Γ' contains an unbalanced switchable block B_0, on which by adequate switching of σ we may assume σ and σ_2 to agree, and an unbalanced σ_2-balanced block B. Without loss of generality we may assume that B_0 and B are joined by a series of doubly balanced blocks $B_1, \ldots, B_k$ (where $k \geq 0$), which by switching we may assume to be all positive in Σ. Let the intersection vertices be $v_i = B_i \cap B_{i+1}$ for $i = 0, \ldots, k$, where $B_{k+1} = B$. Let C be a negative polygon on v_0 in B_0 and D one on v_k in B and let P_i, Q_i be internally disjoint paths in B_i from v_{i-1} to v_i, for $i = 1, \ldots k$. By Lemma 7, $W = CP_1 \ldots P_k D Q_k^{-1} \ldots Q_1^{-1}$ is a face boundary walk in some embedding of Σ. Because W is σ_2-negative, Σ cannot have the σ_2-positive embedding property.

Part (iii), ($\Leftarrow$): Suppose every Γ'_i is switchable or σ_2-balanced. Let Σ be embedded in a surface. Because a face boundary walk W is positive and closed, and a positive closed walk in a switchable or σ_2-balanced doubly signed graph is necessarily σ_2-positive, W is σ_2-positive if it lies within a component of Γ'. Thus if W is σ_2-negative, it must include an isthmus of Γ'.

Suppose a σ_2-negative face boundary walk W does exist. For each component Γ'_i that contains an edge of W, let W'_i be the closed walk obtained from W by deleting every edge not in Γ'_i. Since W is σ_2-negative, some W'_i is σ_2-negative, therefore negative.

We show that there is an isthmus e for which W has the form $eW_1e^{-1}W_2$ with W_1 and W_2 negative. First, an isthmus e in W must be traced once in each direction because W is continuous. Hence W has the form $eW_1e^{-1}W_2$. Now, let T be the tree whose vertices are the subscripts i for which $E(W_i) \neq \emptyset$ and whose edges are the isthmi e of Γ that join Γ'_i and Γ'_j where $i, j \in V(T)$; the endpoints of e in T are i and j. Give to each vertex i of T the sign $\sigma(W'_i)$. We now have a vertex-signed tree in which the number of negative vertices is even and positive. It follows from Lemma 9 that there is an edge e such that each component of $T \backslash e$ has an odd number of negative vertices. This is the edge we wanted, for $\sigma(W_1)$ and $\sigma(W_2)$ are negative because they equal the products of vertex signs in the components of $T \backslash e$.

But since $W = eW_1e^{-1}W_2$, that contradicts Lemma 8. Therefore a σ_2-negative W cannot exist. It follows that Σ has the σ_2-positive embedding property.

Part (iv), (b) $\Rightarrow$ (c): Because an unbalanced closed walk in $(\tilde{\Sigma}, \tilde{\sigma}_2)$ projects to a positive but σ_2-negative closed walk in (Σ, σ_2).

Part (iv), (c) $\Rightarrow$ (d): Without loss of generality we assume Σ connected. Since $(\tilde{\Sigma}, \tilde{\sigma}_2)$ is balanced, Lemma 6 implies that (Σ, σ_2) is switchable or Σ_2 is balanced.

Part (iv), (d) $\Rightarrow$ (b): Trivial.

Part (iv), (a) $\Rightarrow$ (d): Assume (a). By part (i), each block is σ_2-balanced or switchable. If a component of Σ is not itself σ_2-balanced or switchable, it has an unbalanced σ_2-balanced block and an unbalanced switchable block. Let C_1 and C_2 be negative polygons in the former and latter block, respectively, and P a minimal connecting path. Then $C_1 \cup C_2 \cup P$ is a σ_2-negative bias circuit, contradicting the assumption (a).

Part (iv), (b) $\Rightarrow$ (a): Trivial.

Part (v), ($\Rightarrow$): Supposing every lift circuit is σ_2-positive, each block must be σ_2-balanced or switchable, by part (i). If the whole of (Σ, σ_2) is not σ_2-balanced or switchable, it contains a negative polygon C_1 in an unbalanced σ_2-balanced block and a negative polygon C_2 in an unbalanced switchable block, and $C_1 \cup C_2$ is a σ_2-negative lift circuit, contrary to hypothesis.

Part (v), ($\Leftarrow$): All polygons are σ_2-positive if Σ_2 is balanced. If (Σ, σ_2) is switchable, every polygon has the same sign in Σ_2 as in Σ_1. Hence in either case a lift circuit is σ_2-positive whether it consists of one positive polygon or two negative ones. ∎

Chromatic Number

The coloring rule for signed graphs is that one colors the vertices in "colors" $0, \pm 1, \pm 2, \cdots, \pm k$ so that, if vertices x and y are positively adjacent, the colors $c(x)$ and $c(y)$ are unequal, while if x and y are negatively adjacent, $c(x) + c(y) \neq 0$. The (*unrestricted*) *chromatic number* $\chi(\Sigma)$ is the smallest k for which this is possible, while the *zero-free chromatic number* $\chi^*(\Sigma)$ is the smallest k for which it is possible to color without using the color 0. A positive loop precludes colorability; then we say $\chi(\Sigma) = \chi^*(\Sigma) = \infty$. Otherwise, Σ has finite unrestricted and zero-free chromatic numbers. We mention that a negative loop precludes coloring with 0 at the vertex incident to it. It is easy to see that

$$\chi(\Sigma) \leq \chi^*(\Sigma) \leq \chi(\Sigma) + 1.$$

(If $\chi(\Sigma) = k$, the color 0 can be replaced by $+(k+1)$; hence $\chi^*(\Sigma) \leq \chi(\Sigma) + 1$.) This has an interesting consequence. If one studies small chromatic numbers, starting from the smallest possible values $\chi(\Sigma) = \chi^*(\Sigma) = 0$ (which pertain only to the void graph), the zero-free and unrestricted numbers increase alternately: if we write them as an ordered pair (χ, χ^*), all possible pairs have the form (k, k) or $(k, k+1)$; the smallest pair is (0,0), the next smallest is (0,1), then (1,1), (1,2), and so on. Our final theorem tells us that the smallest nontrivial chromatic number pair is (1,1); it corresponds to the class of antibalanced signed graphs; and the next smallest pair (1,2) corresponds to a class that contains both antibalanced and bipartite signed graphs but also others obtained by inserting an antibalanced into a bipartite signed graph. Thus from the standpoint of chromatic number, the closest signed analogs of bipartite graphs are only the antibalanced signed graphs, while the second closest analogs are many more than the antibalanced and bipartite ones to which we were led by evenness properties, even if one restricts attention to block graphs.

Theorem 3. *Let Σ be a signed graph.*

(i) $\chi(\Sigma) = 0 \Longleftrightarrow \Sigma$ *has no edges.*

(ii) $\chi^*(\Sigma) \leq 1 \Longleftrightarrow \Sigma$ *is antibalanced.*

(iii) $\chi(\Sigma) \leq 1 \Longleftrightarrow \Sigma$ *is the union of a signed bipartite graph with bipartition $V = X \cup Y$ (where X or Y may be void) and an antibalanced signed graph with vertex set X.*

Proof. (i) is trivial.

(ii) We need to explain how switching transforms colors. If W is switched, the colors of vertices in W are negated but those of other vertices remain as they were.

If Σ is antibalanced, switch so it is all negative. Color every vertex $+1$. Hence $\chi^*(\Sigma) \leq 1$. Conversely, suppose Σ can be colored using the colors $+1$ and -1. Switching all negatively colored vertices, we see that every edge becomes negative. Thus Σ was antibalanced.

(iii) now follows from Lemma 10. ∎

Lemma 10. *Let $k \geq 0$. A signed graph Σ has chromatic number $\chi(\Sigma) \leq k$ if and only if it consists of a bipartite signed graph with bipartition $V = X \cup Y$ and a signed graph Σ_X on vertex set X satisfying $\chi^*(\Sigma_X) \leq k$.*

Proof. If $\chi(\Sigma) \leq k$, color Σ with the colors $0, \pm 1, \cdots, \pm k$. Take Y to be the set of 0-colored vertices and X its complement. Then Y supports no edges and X induces a subgraph Σ_X which is colored by $\pm 1, \cdots, \pm k$, hence has $\chi^*(\Sigma_X) \leq k$. Conversely, if Σ has the specified form, color X in colors $\pm 1, \cdots, \pm k$ and Y in the color 0. Thus $\chi(\Sigma) \leq k$. ■

References

1. F. Harary, On the notion of balance of a signed graph, Michigan Math. J. 2 (1953-54) 143-146; addendum, ibid., preceding p. 1.

2. F. Harary, On local balance and N-balance in signed graphs, Michigan Math. J. 3 (1955-56) 37-41.

3. S. Lins, Graph-encoded maps, J. Combin. Theory Ser. B 32 (1982) 171-181.

4. S. Lins, Combinatorics of orientation reversing polygons, Aequat. Math. 29 (1985) 123-131.

5. G. Ringel, The combinatorial map color theorem, J. Graph Theory 1 (1977) 141-155.

6. S. Stahl, Generalized embedding schemes, J. Graph Theory 2 (1978) 41-52.

7. K. Truemper, Alpha-balanced graphs and matrices and $GF(3)$-representability of matroids, J. Combin. Theory Ser. B 32 (1982), 112-139.

8. T. Zaslavsky, Biased graphs. I. Bias, balance, and gains; II. The three matroids, J. Combin. Theory Ser. B 47 (1989), 32-52; 51 (1991) 46-72.

9. T. Zaslavsky, Orientation embedding of signed graphs, J. Graph Theory 16 (1992) 399-422.