Discrete Mathematics 39 (1982) 215–228
North-Holland Publishing Company

SIGNED GRAPH COLORING

Thomas ZASLAVSKY

The Ohio State University, 231 West 18th Ave, Columbus, OH 43210, USA

Dedicated to Professor Fred Supnick of the City College of the City University of New York

Received 5 November 1979
Revised August 1981

Coloring a signed graph by signed colors, one has a chromatic polynomial with the same enumerative and algebraic properties as for ordinary graphs. New phenomena are the interpretability only of odd arguments and the existence of a second chromatic polynomial counting zero-free colorings. The generalization to voltage graphs is outlined.

Introduction

A signed graph (a graph whose arcs have been labelled by signs) has, like an ordinary graph, a chromatic polynomial, which appears combinatorially as the function which counts proper colorings of the graph and algebraically essentially as the characteristic polynomial of the associated matroid. In this article we show the existence of colorings and the chromatic polynomial and demonstrate their relationships to acyclic orientations.

The key idea is a simple one: to color a signed graph one needs signed colors. What is remarkable is how closely the development resembles ordinary graph theory. The proper colorings are counted by a polynomial function of the number of colors, a function related to the signed-graphic matroid. Negative arguments give the number of pairs consisting of a coloring and a compatible acyclic orientation. This is exactly as with ordinary graphs. What is different is that one must use an odd number of signed colors, namely $0, \pm 1, \ldots, \pm\mu$, interpreting the chromatic polynomial only at odd arguments, and that there is a parallel coloring theory distinguished merely by not using the color 0, having its own 'balanced' chromatic polynomial which is interpretable only at even arguments. The interaction of these two theories and how it provides algebraic means of computing the chromatic polynomial will appear, together with examples, in [6], the continuation of this article. The generalization to voltage graphs (where an arbitrary group replaces the sign group $\{\pm\}$) we sketch at the end of this paper.

We restate below most of the necessary definitions and results concerning graphs, signs, and orientations. (Proofs appear in [4] and [5], as do a few omitted definitions of lesser importance.) Thus we hope that this article can be read without prior knowledge of signed graphs.

0012-365X/82/0000–0000/$02.75

In conclusion we remark, *pro forma*, that all graphs considered in coloring theory are finite.

1. Signed graphs and signed colors

1.1. *Reminiscences of signed graphs*

We shall consistently denote by Σ a signed graph, by N its node set, and by n the number of nodes.

A *signed graph* Σ consists of an unsigned graph, denoted $|\Sigma|$, whose node and arc sets are denoted $N(\Sigma)$ and $E(\Sigma)$, and a sign function σ which labels each arc (except a half arc) positive or negative. Sometimes we write $\Sigma = (\Gamma, \sigma)$ or (N, E, σ) to indicate that $|\Sigma| = \Gamma$ or that $N(\Sigma) = N$ and $E(\Sigma) = E$. The underlying graph $|\Sigma|$ may have not only multiple arcs and loops but also *half arcs* (which have but one end point, and are not sign-labelled) and *free loops* (which have no endpoints and are required to be labelled $+$). By Σ^* we mean Σ with its half arcs and free loops removed.

An *ordinary* graph is an unsigned graph without half arcs or free loops. For an ordinary graph Γ we employ the following notations:

$\Pi(\Gamma)$ = the set of partitions of the node set into connected blocks,
$\psi_\kappa(\Gamma)$ = the number of ways to color Γ properly using exactly κ unlabelled colors,
$\chi_\Gamma(\lambda)$ = the chromatic polynomial $= \sum \psi_\kappa(\Gamma)(\lambda)_\kappa$, where $(\lambda)_\kappa$ is the falling factorial $\lambda(\lambda-1)\cdots(\lambda-\kappa+1)$.

An arc set S is *balanced* if it contains no half arcs and each circular path in it has positive product of signs. We write $\pi_b(S) = \{W: W$ is the node set of a balanced component of S, and $W \neq \emptyset\}$ and $b(S) = \#\pi_b(S)$. The *rank* of S is given by $\operatorname{rk} S = n - b(S)$. The rank function determines a matroid $G(\Sigma)$ on the arc set; its lattice of *closed sets* or *flats* is denoted $\operatorname{Lat}\Sigma$. The poset of balanced flats is $\operatorname{Lat}^b \Sigma$. $N(S)$ is the set of nodes of arcs of S; $N_u(S)$ is the set of nodes of unbalanced components. A node is *full* if it supports a half arc or negative loop, otherwise *empty*. Σ is *full* if all its nodes are. By $\Sigma^\bullet$ we mean Σ with each empty node filled by an added half arc or negative loop.

Switching Σ by $\nu: N \to \{\pm 1\}$ means reversing the sign of any link whose endpoints have opposite ν values. It does not alter balance.

The *restriction* to an arc set S is the subgraph $\Sigma \mid S = (N, S, \sigma \mid S)$. The *contraction* Σ/S is obtained by switching Σ so every balanced component of S is all positive, coalescing all the nodes of each balanced component, and discarding the remaining nodes and all the arcs in S. Thus $N(\Sigma/S) = \pi_b(S)$. The subgraph *induced* by $X \subseteq N$ is $\Sigma : X = (X, E:X, \sigma \mid E:X)$, where $E:X = \{e \in E: \emptyset \neq N(e) \subseteq X\}$.

The *signed covering graph* of Σ is the unsigned (or, all-positive) graph $\tilde{\Sigma}$ whose node set is $\pm N(\Sigma) = \{\varepsilon v: v \in N(\Sigma),\ \varepsilon \in \{\pm\}\}$ and which has two arcs covering each one of Σ as follows. A link or loop $e:vw$, with sign ε, is covered by two arcs, $\tilde{e}_1: +v, \varepsilon w$ and $\tilde{e}_2: -v, -\varepsilon w$. A free loop is covered by two free loops. A half arc $e:v$ is covered by two half arcs, $e_1: +v$ and $e_2: -v$. The *covering projection* $p:\tilde{\Sigma}\to\Sigma$ is given by $p(\pm v)=v$ and, if $\tilde{e}$ covers e, $p(\tilde{e})=e$. The *opposite* is the involutory automorphism $\tilde{x}\to\tilde{x}^*$ of $\tilde{\Sigma}$ which interchanges the two objects covering each $x\in\Sigma$. A modified $\tilde{\Sigma}$, which has no half arcs yet behaves just like the original $\tilde{\Sigma}$ in respect to closure, orientation, and coloring, covers each half arc $e:v$ by one link $\tilde{e}: +v, -v$. The opposite then fixes $\tilde{e}$

We employ the signed covering (basic or modified) to reduce questions about signed graphs to ones about ordinary graphs. The first relevant fact is this:

Lemma 1.1 [4, Theorem 6.5(i)]. *An arc set $S\subseteq E(\Sigma)$ is closed in $G(\Sigma)\Leftrightarrow p^{-1}(S)$ is closed in $G(\tilde{\Sigma})$.*

An *orientation* τ of Σ (cf. [5]) is obtained by orienting all arcs except free loops. A link or loop, $e:vw$, is oriented by giving to each of its incidences, (v,e) and (w,e), signs $\tau(v,e)$ and $\tau(w,e)$, subject to their product being $-\sigma(e)$. (The two incidences of a loop must be distinguished; their signs may differ.) A half arc $e:v$ is oriented by giving its incidence (v,e) a sign $\tau(v,e)$. The meaning of $\tau(v,e)>0$ is that e points into v, while $\tau(v,e)<0$ means the reverse. Thus a positive arc orients in the usual way. If τ orients Σ, there is a unique *covering orientation* $\tilde{\tau}$ of $\tilde{\Sigma}$ determined by

$$\tau(\varepsilon v,\tilde{e})=\varepsilon\tau(v,p(\tilde{e})).$$

The arcs of $\tilde{\Sigma}$ are oriented by $\tilde{\tau}$ in the way suitable for unsigned (or positive) arcs. A *cycle* of τ is a circuit in $G(\Sigma)$ which, as a subgraph, has neither sources nor sinks. The *cyclic part* of τ, $C(\tau)$, is the union of all cycles.

Lemma 1.2 [5, Theorem 2.1 and Corollary 2.4]. *We have $C(\tilde{\tau})=p^{-1}(C(\tau))$; thus τ is acyclic $\Leftrightarrow\tilde{\tau}$ is. Also $C(\tau)$ is closed in $G(\Sigma)$.*

One tricky point about unsigned graphs must be emphasized. For ordinary graphs the absence of half arcs means that the ordinary theory of graph coloring applies. It is not quite general enough for all unsigned graphs; for coloring purposes they can be regarded as voltage graphs labelled by the trivial group (cf. Section 5), or as having their half arcs replaced by links to an extra node v_0 and their free loops by loops at v_0.

1.2. *Signed colorings*

A coloring of an ordinary graph in λ colors is customarily taken to be a mapping of the nodes into the set $[1,\lambda]=\{1,2,\ldots,\lambda\}$. This convention is unsuitable for signed graphs because it does not admit a definition of proper coloring across a negative arc. For signed graphs one must have signed colors.

Thus if Σ is a signed graph, we define a (*signed*) *coloring of* Σ *in* μ *colors*, or *in* $2\mu+1$ *signed colors*, to be a mapping

$$k : N \to [-\mu, \mu] = \{-\mu, -\mu+1, \ldots, 0, \ldots, \mu-1, \mu\}.$$

A coloring is *zero-free* if it never assumes the value 0. One of the themes of this work is the interplay between colorings and zero-free colorings through the chromatic polynomials.

A second theme is propriety of colorings and their improper sets. An arc e is *improper* for a coloring k if $e:vw$ is a link or loop with endpoints colored $k(w)=\sigma(e)k(v)$ (thus every balanced loop is improper), or $e:v$ is a half arc with endpoint colored $k(v)=0$, or e is a free loop. The set of improper arcs of a coloring k, denoted $I(k)$, may also be called the *set of impropriety* of k. Counting colorings without improper arcs or those with specified improprieties is the heart of signed as ordinary coloring theory. We shall pursue this topic in Section 2. First we must study sets of impropriety and correspondences between proper and improper colorings.

A coloring k of Σ lifts to a unique coloring $\tilde{k}$ of the signed covering graph by the rule

$$\tilde{k}(\varepsilon v) = \varepsilon k(v).$$

The important point is the relationship between the two improper sets. (The proof is a matter of inspection.)

Lemma 1.3. *We have* $I(\tilde{k}) = p^{-1}(I(k))$.

Like a graph, a coloring can be switched; indeed whenever we switch a signed graph Σ we implicitly switch its colorings as well. The mechanism is this: if ν is a switching function on $N(\Sigma)$, we define

$$k^{\nu}(v) = \nu(v)k(v).$$

That is, if a node v is switched ($\nu(v) = -1$), then $k(v)$ is negated to $-k(v)$. The most important thing to notice is that switching preserves sets of impropriety:

$$I(k^{\nu}) = I(k).$$

Switching k to k^{ν} entails switching $\tilde{k}$ also; we call the result $\tilde{k}^{\nu}$.

Lemma 1.4. *The set of impropriety of a coloring k is a closed set of arcs, balanced if k is zero free. For every node v in an unbalanced component of $I(k)$, $k(v)=0$. For v in a balanced component $I(k):X$, $k(v)\neq 0$, except possibly when $I(k):X = E(\Sigma):X$.*

Proof. That $I(k)$ is closed follows from the same property for ordinary graphs because of Lemmas 3 and 1. To illustrate the proof of the description of k, suppose C to be an unbalanced circle in $I(k)$. Say $C = e_1e_2 \cdots e_r$ where $e_i : v_{i-1}v_i$ and $v_0 = v_r$. By the impropriety of C, $k(v_r) = \sigma(e_r)k(v_{r-1}) = \cdots = \sigma(C)k(v_0)$; but $\sigma(C) = -1$, so $k(v_0) = 0$. The other proofs are similar. □

1.3. *Proper colorings*

A coloring is *proper* if its set of impropriety is void; in other words, if

(i) $k(v) \neq \sigma(e)k(w)$ whenever there is an arc $e : vw$;

(ii) $k(v) \neq 0$ whenever there is a half arc at v; and

(iii) Σ has no free loops.

The first constraint implies $k(v) \neq 0$ whenever there is a negative loop at v (take $v = w$). The first and third imply the impossibility of properly coloring a signed graph which has a balanced loop. By Lemma 3, the lifted coloring $\tilde{k}$ is proper if and only if k is proper.

The colorings of Σ can be regarded as proper colorings of the contractions of Σ.

Lemma 1.5. *Let Σ be a signed graph and $\mu \geq 0$ an integer. There is a one-to-one correspondence between all signed colorings of Σ in μ colors and all proper signed colorings of contractions of Σ, in which the zero-free signed colorings correspond to the zero-free proper colorings of contractions by balanced sets.*

A coloring k of Σ corresponds to the proper coloring k' of $\Sigma/I(k)$ determined by first switching Σ (and k) until every balanced component of $I(k)$ is positive, then defining $k'(B) = k(v)$ for $v \in B \in \pi_b(I(k))$.

A proper coloring k' of Σ/A corresponds to the coloring k of Σ determined by first switching Σ until every balanced component of A is positive, then defining $k \mid B = k'(B)$ for each $B \in \pi_b(A)$ and $k \mid N_u(A) = 0$, then reversing the switching of Σ and k.

Note that the correspondence, to be well defined, requires that the switching scheme used on Σ in the first part (where k is given) depend only on the set $I(k)$, while that in the second part depend only on A, not k'. But given this, it does not matter which switching scheme is chosen for each set.

The proof is straightforward and is omitted.

2. Counting the coloring ways

2.1. *Polynomials*

The *chromatic polynomial* $\chi_\Sigma(\lambda)$ of a signed graph Σ is the function defined for odd positive arguments $\lambda = 2\mu + 1$ whose value equals the number of proper signed colorings of Σ in μ colors. The *balanced chromatic polynomial* $\chi_\Sigma^b(\lambda)$, defined for even positive arguments $\lambda = 2\mu$, is the function which counts zero-free proper signed colorings in μ colors. That both functions are polynomials will be proved shortly.

The *Whitney polynomial* of Σ is the generating function of all colorings, classified by the rank of the set of impropriety:

$$w_\Sigma(x, 2\mu+1) = \sum_k x^{\operatorname{rk} I(k)},$$

summed over all signed colorings in μ colors. The coefficient of x^r, the number of colorings whose set of impropriety has rank r, may equally be regarded as the number of proper signed colorings in μ colors of contractions Σ/A having $n-r$ nodes—this follows from Lemma 1.5. The *balanced Whitney polynomial*,

$$w_\Sigma^b(x, 2\mu) = \sum_{k^*} x^{\operatorname{rk} I(k^*)},$$

is summed over zero-free colorings only. Its coefficient of x^r is both the number of zero-free colorings k^* whose improper set has rank r, and the number of zero-free proper signed colorings of $(n-r)$-node contractions Σ/A where A is balanced.

By interpreting the Whitney polynomials in terms of contractions of Σ we see that

$$w_\Sigma(x, \lambda) = \sum_{A \subseteq E(\Sigma)} x^{\operatorname{rk} A} \chi_{\Sigma/A}(\lambda), \tag{2.1}$$

$$w_\Sigma^b(x, \lambda) = \sum_{\substack{A \subseteq E(\Sigma) \\ \text{balanced}}} x^{\operatorname{rk} A} \chi_{\Sigma/A}^b(\lambda), \tag{2.2}$$

if λ is positive and is odd in the first equation, even in the second. (It is sufficient to sum over closed A, since otherwise the contribution of Σ/A is 0.) Since w_Σ and w_Σ^b are polynomials—a consequence of (2.1), (2.2), and Theorem 2—these identities are valid for all numbers λ.

Balanced and ordinary graphs. If $\Sigma = (\Gamma, \sigma)$ is balanced we can assume, switching as necessary, that it is all-positive. Then a signed coloring k has the same set of impropriety whether regarded as a coloring of Σ or of the underlying graph Γ. Thus we see that a balanced signed graph and its underlying unsigned graph have the same chromatic and Whitney polynomials (and it is easy to see that the balanced polynomials equal the unbalanced ones as well). Note that the Whitney polynomial of an unsigned graph is defined like that of a signed graph and satisfies an identity analogous to (2.1).

2.2. *Exact coloring*

To prove that the chromatic and balanced chromatic functions are polynomials we shall study *exact* colorings: signed colorings of Σ in μ colors which use all the magnitudes $1, 2, \ldots, \mu$. (Whether 0 is used or not is immaterial.)

Rather than colorings themselves we shall count symmetry classes. A symmetry class may be regarded as an unlabelled signed coloring. A coloring k using the signed colors $0, \pm 1, \ldots, \pm\mu$ is acted on by the permutations $\alpha \in \mathfrak{S}_\mu$ and the reflections ρ_i, $1 \leqslant i \leqslant \mu$, in the following manner: α permutes the color magnitudes, while ρ_i reverses the sign of every color whose magnitude is i. Thus the hyperoctahedral group $\mathfrak{O}_\mu$, generated by $\mathfrak{S}_\mu$ and the reflections, acts as a symmetry group on colorings. It is easy to check that $I(k^\gamma) = I(k)$ for any $\gamma \in \mathfrak{O}_\mu$; thus propriety is preserved. Furthermore, if k is exact, then $k^\gamma = k$ only if γ is the identity. Hence exact colorings come in symmetry classes of size $2^\mu \mu!$, the order of $\mathfrak{O}_\mu$. Let

$\psi_\mu(\Sigma)$ = the number of symmetry classes of exact proper signed colorings of Σ in μ colors,

$\psi^*_\mu(\Sigma)$ = the number which are zero-free.

Clearly there is only one class using $\mu = n$ colors and it is zero-free, thus $\psi_n(\Sigma) = \psi^*_n(\Sigma) = 1$; and there are none using more than n colors.

We shall compute χ_Σ and $\chi^\flat_\Sigma$ in terms of the $\psi_\mu(\Sigma)$ and $\psi^*_\mu(\Sigma)$, and vice versa. We need the double falling factorial:

$$_2(\lambda)_r = \lambda(\lambda - 2) \cdots (\lambda - 2r + 2).$$

Lemma 2.1. *We have (if $\lambda \geqslant 0$ and is odd for χ_Σ, even for $\chi^\flat_\Sigma$):*

$$\chi_\Sigma(\lambda) = \sum_{\kappa=0}^{n} \psi_\kappa(\Sigma) \cdot {}_2(\lambda - 1)_\kappa,$$

$$\chi^\flat_\Sigma(\lambda) = \sum_{\kappa=0}^{n} \psi^*_\kappa(\Sigma) \cdot {}_2(\lambda)_\kappa.$$

Conversely if $\mu \geqslant 0$ is integral:

$$\psi_\mu(\Sigma) = \frac{(-1)^\mu}{2^\mu \mu!} \sum_{\kappa=0}^{\mu} (-1)^\kappa \binom{\mu}{\kappa} \chi_\Sigma(2\kappa + 1)$$

$$\psi^*_\mu(\Sigma) = \frac{(-1)^\mu}{2^\mu \mu!} \sum_{\kappa=0}^{\mu} (-1)^\kappa \binom{\mu}{\kappa} \chi^\flat_\Sigma(2\kappa)$$

To prove this we can imitate the method used for ordinary graphs (cf. [1]). We omit the details.

2.3. *Properties of the chromatic polynomials*

In perfect analogy to ordinary chromatic theory (as in [1]) we have from Lemma 1:

Theorem 2.2. *If Σ is a signed graph with n nodes, then $\chi_\Sigma(\lambda)$ and $\chi^\flat_\Sigma(\lambda)$ are polynomial functions of λ, monic of degree n.*

Now the chromatic polynomials can be evaluated for all numbers λ and the parenthetical reservation in Lemma 1 can be ignored.

Continuing the analogy with ordinary chromatic theory is the behavior of the signed-graphic polynomials under deletion and contraction. It can be proved by counting proper colorings.

Theorem 2.3. *Let Σ be a signed graph and $e \in E(\Sigma)$. Then*

$$\chi_\Sigma(\lambda) = \chi_{\Sigma \setminus e}(\lambda) - \chi_{\Sigma/e}(\lambda)$$

and, if e is not a half arc or negative loop,

$$\chi_\Sigma^b(\lambda) = \chi_{\Sigma \setminus e}^b(\lambda) - \chi_{\Sigma/e}^b(\lambda).$$

On the coloring interpretation. We have seen that $\chi_\Sigma(\lambda)$ for a positive integer λ is the number of colorings of Σ in λ colors, but only when λ is odd. To suggest why more cannot be expected let us consider the signed graph $\pm K_2^\circ$, consisting of two nodes linked by one positive and one negative arc, with a negative loop at each node. Since $\psi_0 = \psi_1 = 0$, Lemma 1 yields

$$\chi_{\pm K_2^\circ}(\lambda) = (\lambda - 1)(\lambda - 3).$$

We see that $\chi_{\pm K_2^\circ}(2) = -1$, demolishing the hope that $\chi_\Sigma(\lambda)$ counts colorings for all positive integral λ. Just what $\chi_\Sigma(\lambda)$ can mean for even arguments is a delicate question, not to be considered here.

2.4. *The algebraic connection*

As in ordinary graph theory there is an alternative algebraic definition of the chromatic polynomial. That and the connection with the matroid $G(\Sigma)$ are the content of the following theorem. Let μ be the Möbius function of Lat Σ (cf. [2]) and let $p(\lambda)$ be the characteristic polynomial of $G(\Sigma)$. (We remind the reader of the convention that $\mu(\emptyset, A) = 0$ if the null set is not closed.)

Theorem 2.4. *The chromatic polynomial of a signed graph Σ satisfies the equation*

$$\begin{aligned} \chi_\Sigma(\lambda) &= \sum_{S \subseteq E(\Sigma)} \lambda^{b(S)} (-1)^{\#S} \\ &= \sum_{A \in \operatorname{Lat} \Sigma} \mu(\emptyset, A) \lambda^{b(A)} = \lambda^{b(\Sigma)} p(\lambda). \end{aligned} \tag{2.3}$$

The balanced chromatic polynomial satisfies

$$\chi_\Sigma^b(\lambda) = \sum_{\substack{S \subseteq E(\Sigma) \\ \text{balanced}}} \lambda^{b(S)} (-1)^{\#S} = \sum_{A \in \operatorname{Lat}^b \Sigma} \mu(\emptyset, A) \lambda^{b(A)}. \tag{2.4}$$

These equations show that the chromatic and Whitney polynomials are invariants of switching classes, because they depend only on the matroid, which depends only on the switching class of Σ (by [4, Corollary 5.4]).

Proof. The subset and flat versions of each formula are equivalent by a standard expansion of the Möbius function. We prove the balanced Möbius formula (2.4) by Möbius inversion (cf. [2]); (2.3) is similar. That $\chi(\lambda)$ and $p(\lambda)$ are related follows from the known Möbius expansion of $p(\lambda)$.

By Lemma 1.5, the number of zero-free signed colorings of Σ using κ colors whose set of impropriety equals S is $\chi^{\mathrm{b}}_{\Sigma/S}(2\kappa)$. Since the set of impropriety is closed and balanced (Lemma 1.4),

$$\sum_{S\in \mathrm{Lat}^{\mathrm{b}}\Sigma} \chi^{\mathrm{b}}_{\Sigma/S}(2\kappa) = (2\kappa)^n.$$

In order to invert we must have a similar identity for every contraction Σ/A where A is a balanced flat. We get it by replacing Σ by Σ/A; thus

$$\sum_{S\geqslant A} \chi^{\mathrm{b}}_{\Sigma/S}(2\kappa) = (2\kappa)^{b(A)},$$

summed over flats S for which S/A is balanced in Σ/A. By Lemma 5 below that means S is really varying over all $S\in \mathrm{Lat}^{\mathrm{b}}\Sigma$ such that $S\geqslant A$. Inverting and setting $A=\emptyset$ yields (2.4)—except in the trivial case when $\emptyset$ is not closed. □

Lemma 2.5 [4, Lemma 4.1]. *Suppose A is balanced in Σ and $S\supseteq A$. Then S is balanced in $\Sigma \Leftrightarrow S/A$ is balanced in Σ/A.*

3. Pairs of colorings and orientations

Following the pioneering work of Stanley [3] we can interpret the chromatic polynomial evaluated at negative as well as positive odd integers. Let us first observe that a signed coloring k determines an orientation of each proper arc if we require that, for a half arc $e:v$,

$$\tau(v, e)k(v)>0,$$

while for an arc $e:vw$,

$$\tau(v, e)k(v)+\tau(w, e)k(w)>0.$$

Since the left side of the latter equals $\tau(v, e)[k(v)-\sigma(e)k(w)]$, which is non-zero precisely because k is proper on e, there is a unique orientation of the desired kind. If k was proper, we have now oriented the entire signed graph. Note that the lifted coloring $\tilde{k}$ determines by this rule the lifted orientation $\tilde{\tau}$.

Proper pairs. An orientation τ and a coloring k are a *proper pair* if, for any arc $e:vw$,

$$\tau(v,e)k(v)+\tau(w,e)k(w)>0,$$

and for any half arc $e:v$,

$$\tau(v,e)k(v)>0.$$

One can easily calculate that $(\tilde{k},\tilde{\tau})$ is a proper pair if and only if (k,τ) is.

Theorem 3.1. *If k is a proper signed coloring of the signed graph Σ, there is a unique orientation τ of Σ such that (k,τ) is a proper pair; and τ is acyclic. If k is not proper, it is in no proper pair. The number of proper pairs allowing μ colors is $\chi_\Sigma(2\mu+1)$. The number involving only zero-free colorings is $\chi_\Sigma^b(2\mu)$.*

The only new statement, that τ is acyclic, follows from Theorem 3 below.

Compatible pairs. A coloring k and an orientation τ are *compatible* if

$$\tau(v,e)k(v)+\tau(w,e)k(w)\geq 0$$

whenever $e:vw$ is a whole arc and

$$\tau(v,e)k(v)\geq 0$$

whenever $e:v$ is a half arc. (If $e:vv$ is a negative loop the first constraint entails $\tau(v,e)k(v)\geq 0$. Thus a half arc is equivalent to a negative loop.) Compatibility, like propriety, is preserved by switching. The lift $(\tilde{k},\tilde{\tau})$ is compatible if and only if (k,τ) is.

A proper pair is just a compatible pair involving a proper coloring. We have counted proper pairs; now we wish to count compatible pairs. The starting point is a lemma which shows how compatibility of pairs, propriety of colorings, and cyclicity of orientations are related to each other.

Lemma 3.2. *Let k be a coloring and τ an orientation of Σ, let $S\subseteq I(k)$, and let k' and τ' be the induced coloring and orientation of Σ/S. Then (k,τ) is compatible $\Leftrightarrow(k',\tau')$ is. In case $S=I(k)$, (k,τ) is compatible $\Leftrightarrow(k',\tau')$ is proper.*

Proof. The definition of k' on Σ/S is like that on $\Sigma/I(k)$ given in Lemma 1.5. The proof depends on the switching invariance of compatibility. Thus one can switch, then verify that each arc is equally compatible in Σ and Σ/S. As for the case $S=I(k)$, since k' is proper (by Lemma 1.5), (k',τ') is proper if it is compatible. □

Theorem 3.3. *If (k,τ) is a compatible pair, then $C(\tau)\leq I(k)$, τ induces an acyclic orientation τ' on $\Sigma/I(k)$, and if k' is the induced coloring of $\Sigma/I(k)$, then (k',τ') is a proper pair.*

Proof. We lift the problem to the signed covering graph. For an ordinary graph like $\tilde{\Sigma}$ it is easy to prove, by tracing inequalities around a cycle of $\tilde{\tau}$, that $\tilde{k}$ is constant on each cycle. Hence $I(\tilde{k}) \supseteq C(\tilde{\tau})$. Then $I(k) \supseteq C(\tau)$ by Lemmas 1.2 and 1.3.

If (k, τ) is proper therefore τ is acyclic. But by Lemma 2, (k', τ') is a proper pair. That completes the proof. □

Lemma 3.4. *Let e be an arc of Σ, k a coloring, and k′ the induced coloring of Σ/e, if there is one. If τ is an orientation of Σ, let τ_e be τ with e reversed.*

If (k, τ) and (k, τ_e) are both compatible, then k is improper on e and both $(k, \tau\backslash e)$ and $(k', \tau/e)$ are compatible pairs in $\Sigma\backslash e$ and Σ/e, respectively.

If (k, τ) is compatible but (k, τ_e) is not, then $(k, \tau\backslash e)$ is compatible but k is proper on e so does not induce a coloring of Σ/e.

If neither (k, τ) nor (k, τ_e) is compatible, then $(k, \tau\backslash e)$ is incompatible and, if there is a k′, $(k', \tau/e)$ is incompatible.

Proof. In the first case, k is clearly improper on e. The rest is from Lemma 2. The second case is obvious. In the third case, clearly $(k, \tau\backslash e)$ must be incompatible. If k' exists, it does so because $e \in I(k)$. Then $(k', \tau/e)$ is incompatible by Lemma 2. □

Theorem 3.5. *Let Σ be a signed graph and μ a nonnegative integer. The number of compatible pairs (k, τ) in which τ is an acyclic orientation and k is a signed coloring in μ colors is $(-1)^n\chi_\Sigma(-(2\mu+1))$. The number in which k is also zero-free is $(-1)^n\chi_\Sigma^b(-2\mu)$.*

Proof. All the compatible pairs in Σ correspond to all those of $\Sigma\backslash e$ and Σ/e (by Lemma 4; the correspondence is clearly bijective), as shown in the following table:

Compatible in Σ		Compatible in $\Sigma\backslash e$, Σ/e
(k, τ) and (k, τ_e)	$\to$	$(k, \tau\backslash e)$ and $(k', \tau/e)$,
(k, τ), not (k, τ_e)	$\to$	$(k, \tau\backslash e)$, not $(k', \tau/e)$,
neither		neither

We can deduce from Theorem 3.1 of [5] that in the first line the same number of pairs on each side involve acyclic orientations. In the second line $\tau\backslash e$ is acyclic if τ is; conversely we must show that, if $\tau\backslash e$ is acyclic, τ has to be too.

Suppose that (k, τ) is compatible but (k, τ_e) is not, and that τ_e is acyclic but τ is not. The compatibilities require (if $e: vw$ is a link or loop; the half-arc case is similar)

$$\tau(v, e)k(v) + \tau(w, e)k(w) > 0.$$

That entails k is proper on e. But on the other hand the acyclicities require $e \in C(\tau)$, whence $e \in I(k)$ by Theorem 3. But this is a contradiction. Thus the troublesome case cannot exist.

Let $P(\Sigma)$ = the set of all compatible pairs (k, τ) in which τ is acyclic. We have shown that $P(\Sigma)$ is in one-to-one correspondence with the disjoint union $P(\Sigma \setminus e) \cup P(\Sigma / e)$. Now a standard inductive argument using Theorem 2.3 proves the general case of the theorem.

If we restrict to zero-free pairs we still have the same one-to-one correspondence, so long as e is a link Thus this case too follows by induction. □

Colored and oriented contractions. From Theorem 3.5 we immediately get interpretations of the Whitney polynomial and the balanced Whitney polynomial.

Corollary 3.6. *The number of compatible triples $(k, \tau, \Sigma/A)$, where Σ/A is a c-node contraction of Σ, τ is an acyclic orientation of Σ/A, and k is a signed coloring of Σ/A in μ colors which is compatible with τ, equals the coefficient of x^{n-c} in $(-1)^n w_\Sigma(-x, -(2\mu+1))$. The number involving only zero-free colorings equals the coefficient of x^{n-c} in $(-1)^n w_\Sigma^b(-x, -2\mu)$.*

4. Orientations, hyperplanes, and the acyclotope

The number of acyclic orientations of a signed graph and its contractions is an immediate consequence of Theorem 3.5 and Corollary 3.6. We simply set $\mu = 0$.

Corollary 4.1. *The number of acyclic orientations of Σ, $o(\Sigma)$, is given by*

$$o(\Sigma) = (-1)^n \chi_\Sigma(-1). \tag{4.1}$$

If $o_k(\Sigma)$ denotes the number of acyclic orientations of all contraction graphs Σ/A which have k nodes, then

$$\sum_{k=0}^{n} o_k(\Sigma) x^{n-k} = (-1)^n w_\Sigma(-x, -1). \tag{4.2}$$

Alternative interpretations of Corollary 1 follow from the representations of Σ in $\mathbb{R}^n$ as an arrangement of hyperplanes $H[\Sigma]$ and as a zonotope $Z[\Sigma]$, the *acyclotope* of Σ. From [5] we recall the definitions. Let the standard basis of $\mathbb{R}^n$ be indexed by the nodes, so it is $\{b_v : v \in N\}$. Each arc $e: vw$ determines a line segment $S_e = \operatorname{conv}\{\pm(b_v - \sigma(e) b_w)\}$; a half arc $e: v$ determines $S_e = \operatorname{conv}\{\pm b_v\}$; and a free loop $e: \emptyset$ determines $S_e = \{0\}$. Then the definitions are

$$H[\Sigma] = \{S_e^\perp : e \in E(\Sigma)\}, \qquad Z[\Sigma] = \sum_e S_e.$$

By Theorem 4.2 and Corollary 4.5 of [5] the regions and, more generally, the k-dimensional faces of $H[\Sigma]$ are in one-to-one correspondence respectively with the acyclic orientations of Σ and of the k-node contractions. Thus we can reformulate Corollary 1 as a statement about $H[\Sigma]$.

Corollary 4.1′. *Corollary 1 remains true if $o(\Sigma)$ is interpreted as the number of regions of $H[\Sigma]$ and $o_k(\Sigma)$ as the number of k-dimensional faces.*

Since the $(n-k)$-faces of the acyclotope $Z[\Sigma]$ correspond to the k-faces of $H[\Sigma]$, we also have a statement about the faces of $Z[\Sigma]$.

Corollary 4.1″. *Corollary 1 remains true if $o(\Sigma)$ is interpreted as the number of vertices of $Z[\Sigma]$ and $o_k(\Sigma)$ as the number of $(n-k)$-dimensional faces.*

In [6], the continuation of this article, we shall find explicit expressions for $o(\Sigma)$ and $o_k(\Sigma)$ in a number of interesting examples.

5. Appendix on voltage graph coloring

A *voltage graph* Φ consists of an unsigned graph $\Gamma=(N,E)$ together with a *voltage group* $\mathfrak{G}$ and a *voltage*, a mapping $\varphi: E^* \to \mathfrak{G}$. We assume that $\varphi(e^{-1})=\varphi(e)^{-1}$, where e^{-1} denotes the arc e taken in the opposite direction, and that $\varphi \mid \{\text{free loops}\} \equiv 1$.

Since the notion of balance applies to voltage graphs one can define both ordinary and balanced chromatic polynomials by the algebraic formulas of Theorem 2.4. There is also a concrete interpretation by means of $\mathfrak{G}$-labelled colorings of Φ. Say $\mathscr{G}$ has order m. The *color set* with μ colors (or $m\mu+1$ labelled colors) is

$$K_\mu=(\{0,1,\ldots,\mu\}\times\mathfrak{G})/(\{0\}\times\mathfrak{G});$$

in other words we regard all $(0, g)$ as the same color. In an obvious way $\mathfrak{G}$ acts from the right on k_μ. A *$\mathfrak{G}$-coloring of Φ in μ colors* is a mapping $k: N \to K_\mu$; it is *proper* if, for each arc $e: v \to w$, $k(w) \neq k(v)\varphi(e)$, and for each half arc $e: v$, $k(v) \neq (0, g)$, and there are no free loops. The function

$$\chi_\Phi(m\mu+1)=\text{the number of proper } \mathfrak{G}\text{-colorings of } \Phi \text{ in } \mu \text{ colors}$$

equals the algebraically defined chromatic polynomial;

$$\chi_\Phi^b(m\mu)=\text{the number of zero-free proper colorings}$$

equals the algebraic balanced polynomial. The proofs are exactly as for signed graphs. Moreover the two polynomials are related just as are those of signed graphs (see the balanced expansion formula in [6]). The big difference is that voltage graphs do not in general have orientations. Thus only positive arguments have combinatorial interpretations.

References

[1] R.C. Read, An introduction to chromatic polynomials, J. Combin. Theory 4 (1968) 52–71. MR 37 #104.
[2] G.-C. Rota, On the foundations of combinatorial theory, I. Theory of Möbius functions, Z. Wahrscheinlichkeitstheorie 2 (1964) 340–368. MR 30 #4688.
[3] R.P. Stanley, Acyclic orientations of graphs, Discrete Math. 5 (1973) 171–178. MR 47 #6537.
[4] T. Zaslavsky, Signed graphs, Discrete Appl. Math. 4 (1982), to appear.
[5] T. Zaslavsky, Orientation of signed graphs, submitted.
[6] T. Zaslavsky, Chromatic invariants of signed graphs, Discr. Math., to appear.