

AN ELEMENTARY CHROMATIC REDUCTION FOR GAIN GRAPHS AND SPECIAL HYPERPLANE ARRANGEMENTS¹

Pascal Berthomé²

Laboratoire d'Informatique Fondamentale d'Orléans,
ENSI de Bourges, Université d'Orléans
88, Bld Lahitolle
18020 Bourges Cedex, France
`pascal.berthome@ensi-bourges.fr`

Raul Cordovil³

Departamento de Matematica
Instituto Superior Técnico
Av. Rovisco Pais
1049-001 Lisboa, Portugal
`rcordov@math.ist.utl.pt`

David Forge

Laboratoire de Recherche en Informatique, UMR 8623
Bât. 490, Université Paris-Sud
91405 Orsay Cedex, France
`forge@lri.fr`

Véronique Ventos

INRIA Futurs-Projet GEMO,
Parc Club Orsay Université
4, rue Jacques Monod – Bât. G
91405 Orsay Cedex, France
`ventos@lri.fr`

Thomas Zaslavsky⁴

Department of Mathematical Sciences
Binghamton University (SUNY)
Binghamton, NY 13902-6000, U.S.A.
`zaslav@math.binghamton.edu`

¹October 2, 2007.

Mathematics Subject Classifications (2010): *Primary* 05C22; *Secondary* 05C15, 52C35.

Key words: Gain graph, integral gain graph, deletion-contraction, loop nullity, weak chromatic function, invariance under simplification, total chromatic polynomial, integral chromatic function, modular chromatic function, affinographic hyperplane, Catalan arrangement, Linial arrangement, Shi arrangement.

²Work performed while at the Laboratoire de Recherche en Informatique, Université Paris-Sud.

³Partially supported by FCT (Portugal) through the program POCTI. Much of Cordovil's work was performed while visiting the LRI, Université Paris-Sud.

⁴Much of Zaslavsky's work was performed while visiting the LRI, Université Paris-Sud.

Abstract

A gain graph is a graph whose edges are labelled invertibly by *gains* from a group. *Switching* is a transformation of gain graphs that generalizes conjugation in a group. A *weak chromatic function* of gain graphs with gains in a fixed group satisfies three laws: deletion-contraction for links with neutral gain, invariance under switching, and nullity on graphs with a neutral loop. The laws are analogous to those of the chromatic polynomial of an ordinary graph, though they are different from those usually assumed of gain graphs or matroids. The three laws lead to the *weak chromatic group* of gain graphs, which is the universal domain for weak chromatic functions. We find expressions, valid in that group, for a gain graph in terms of minors without neutral-gain edges, or with added complete neutral-gain subgraphs, that generalize the expression of an ordinary chromatic polynomial in terms of monomials or falling factorials. These expressions imply relations for all switching-invariant functions of gain graphs, such as chromatic polynomials, that satisfy the deletion-contraction identity for neutral links and are zero on graphs with neutral loops. Examples are the total chromatic polynomial of any gain graph, including its specialization the zero-free chromatic polynomial, and the integral and modular chromatic functions of an integral gain graph.

We apply our relations to some special integral gain graphs including those that correspond to the Shi, Linial, and Catalan arrangements, thereby obtaining new evaluations of and new ways to calculate the zero-free chromatic polynomial and the integral and modular chromatic functions of these gain graphs, hence the characteristic polynomials and hypercubical lattice-point counting functions of the arrangements. The proof involves gain graphs between the Catalan and Shi graphs whose polynomials are expressed in terms of descending-path vertex partitions of the graph of (-1) -gain edges.

We also calculate the total chromatic polynomial of any gain graph and especially of the Catalan, Shi, and Linial gain graphs.

1 Introduction

To calculate the chromatic polynomial $\chi_\Gamma(q)$ of a simple graph there is a standard method that comes in two forms. One can delete and contract edges, repeatedly applying the identity $\chi_\Gamma = \chi_{\Gamma \setminus e} - \chi_{\Gamma/e}$ and the reduction $\chi_\Gamma(q) = 0$ if Γ has a loop, to reduce the number of edges to zero. One ends up with a weighted sum of chromatic polynomials of edgeless graphs, i.e., of monomials q^k . Or, one can add missing edges using the opposite identity, $\chi_\Gamma = \chi_{\Gamma \cup e} + \chi_{\Gamma/e}$, until χ_Γ becomes a sum of polynomials of complete graphs, i.e., of falling factorials $(q)_k$. We extend these approaches to gain graphs, where the edges are orientably labelled by elements of a group. The resulting formulas are slightly more complex than those for simple graphs, but they can be used to compute examples; we show this with gain graphs related to the Shi, Linial, and Catalan hyperplane arrangements.

In a *gain graph*, having edges labelled orientably means that reversing the direction of an edge inverts the label (the *gain* of the edge). Gain graphs, like ordinary graphs,

have chromatic polynomials, which for various choices of gain group give the characteristic polynomials of interesting arrangements of hyperplanes. In particular, when the gain group is \mathbb{Z} , the additive group of integers, the edges correspond to *integral affinographic hyperplanes*, that is, hyperplanes of the form $x_j = x_i + m$ for integers m . Arrangements (finite sets) of hyperplanes of this type, which include the Shi arrangement, the Linial arrangement, and the Catalan arrangement, have gained much interest in recent years. The fact that the chromatic polynomials of gain graphs satisfy the classic deletion-contraction reduction formula has important consequences, e.g., a closed-form formula, and also a method of computation that can significantly simplify computing the polynomials of the Shi, Linial, and Catalan arrangements.

The deletion-contraction relation for functions can be viewed abstractly at the level of a Tutte group, which means that we take the free abelian group $\mathbb{Z}\mathcal{G}(\mathfrak{G})$ generated by all gain graphs with fixed gain group \mathfrak{G} , and from deletion-contraction we infer algebraic relations satisfied in a quotient of $\mathbb{Z}\mathcal{G}(\mathfrak{G})$. This gives a group we call the *neutral chromatic group* of gain graphs; it is a special type of Tutte group. In particular, the relations reduce the number of generators. These relations on graphs then, by functional duality, automatically generate the original deletion-contraction relations on chromatic polynomials and other functions of a similar type that we call *weak chromatic functions* or, if they are invariant under switching (defined in the next section), *simplification*, and isomorphism, *weak chromatic invariants*.

Our investigation in [3] of the number of integer lattice points in a hypercube that avoid all the hyperplanes of an affinographic hyperplane arrangement led us to functions of integral gain graphs, the integral and modular chromatic functions, that count proper colorations of the gain graph from the color set $\{1, \dots, q\}$, treated as either integers or modulo q . Like the chromatic polynomial, they satisfy a deletion-contraction identity, but (in the case of the integral chromatic function) only for links with neutral gain. That fact is part of what suggested our approach, and evaluating these chromatic functions for the Catalan, Shi, and Linial arrangements is part of our main results.

A brief outline: The first half of the paper develops the general theory of functions on gain graphs that satisfy deletion-contraction for neutral links and are zero on gain graphs with neutral loops, in terms of the neutral chromatic group. It also develops, first of all, the corresponding theory for ordinary graphs, since that is one way we prove the gain-graphic reduction formulas. The second half applies the theory to the computation of chromatic functions of the Catalan, Shi, and Linial gain graphs and a family of graphs intermediate between the Catalan and Shi graphs. The latter can be computed in terms of partitions into descending paths of the vertex set of a graph. This half also shows how to compute the total chromatic polynomial in terms of the zero-free chromatic polynomial; this in particular gives the chromatic polynomials of the Catalan, Shi, and Linial (and intermediate) graphs, although since the results are not as interesting as the method we do not state them.

2 Basic definitions

For a nonnegative integer k , $[k]$ denotes the set $\{1, 2, \dots, k\}$, the empty set if $k = 0$. The set of all partitions of $[n]$ is Π_n . The falling factorial is $(x)_m = x(x-1)\cdots(x-m+1)$.

Suppose $\mathcal{P}(E)$ is the power set of a finite set E , $S \mapsto \bar{S}$ is a closure operator on E , \mathcal{F} is the class of closed sets, and μ is the Möbius function of \mathcal{F} . If the empty set is not closed, then $\mu(\emptyset, A)$ is defined to be 0 for all $A \in \mathcal{F}$. It is known that:

Lemma 2.1. *For each closed set A , $\mu(\emptyset, A) = \sum_S (-1)^{|S|}$, summed over edge sets S whose closure is A .*

Proof. If \emptyset is closed, this is [1, Prop. 4.29]. Otherwise, $\sum_{\bar{S}=A} (-1)^{|S|} = \sum_{S \subseteq \emptyset} (-1)^{|S|} x = 0x = 0$, where x is a sum over some subsets of $E \setminus \emptyset$. \square

Our usual name for a graph is $\Gamma = (V, E)$. Its vertex set is $V = \{v_1, v_2, \dots, v_n\}$. All our graphs are finite. Edges in a graph are of two kinds: a *link* has two distinct endpoints; a *loop* has two coinciding endpoints. Multiple edges are permitted. Edges that have the same endpoints are called *parallel*. The *simplification* of a graph is obtained by removing all but one of each set of parallel edges, including parallel loops. (This differs from the usual definition, in which loops are deleted also.) The complement of a graph Γ is written Γ^c ; this is the simple graph whose adjacencies are complementary to those of Γ . The complete graph with W as its vertex set is K_W . For a partition π of V , K_π denotes a complete graph whose vertex set is π .

If $S \subseteq E$, we denote by $c(S)$ the number of connected components of the spanning subgraph (V, S) (which we usually simply call the “components of S ”) and by $\pi(S)$ the partition of V into the vertex sets of the various components. The complement of S is $S^c = E \setminus S$.

An edge set S in Γ is *closed* if every edge whose endpoints are joined by a path in S is itself in S . $\mathcal{F}(\Gamma)$ denotes the lattice of closed sets of Γ .

Contracting Γ by a partition π of V means replacing the vertex set by π and changing the endpoints v, w of an edge e to the blocks $B_v, B_w \in \pi$ that contain v and w (they may be equal); we write Γ/π for the resulting graph.

A vertex set is *stable* if no edge has both endpoints in it. A *stable partition* of Γ is a partition of V into stable sets; let $\Pi^*(\Gamma)$ be the set of all such partitions. Contracting a graph by a stable partition π means collapsing each block of π to a vertex and then simplifying parallel edges; there will be no loops.

A *gain graph* $\Phi = (\Gamma, \varphi)$ consists of an underlying graph Γ and a function $\varphi : E \rightarrow \mathfrak{G}$ (where \mathfrak{G} is a group), which is orientable, so that if e denotes an edge oriented in one direction and e^{-1} the same edge with the opposite orientation, then $\varphi(e^{-1}) = \varphi(e)^{-1}$. The group \mathfrak{G} is called the *gain group* and φ is called the *gain mapping*. A *neutral edge* is an edge whose gain is the neutral element of the group, that is, $1_{\mathfrak{G}}$, or for additive groups 0. The *neutral subgraph* of Φ is the subgraph $\Gamma_0 := (V, E_0)$ where E_0 is the set of neutral edges.

We sometimes use the simplified notations e_{ij} for an edge with endpoints v_i and v_j , oriented from v_i to v_j , and ge_{ij} for such an edge with gain g ; that is, $\varphi(ge_{ij}) = g$. (Thus ge_{ij} is the same edge as $g^{-1}e_{ji}$.)

A second way to describe a gain graph, equivalent to the definition, is as an ordinary graph Γ , having a set of gains for each oriented edge e_{ij} , in such a way that the gains of e_{ji} are the inverses of those of e_{ij} . For instance, K_n with additive gains 1, -1 on every edge is a gain graph that has edges $1e_{ij}$, and $-1e_{ij}$ for every $i \neq j$. Since this is a gain graph, $1e_{ij}$ and $-1e_{ji}$ are the same edge.

Two gain graphs are *isomorphic* if there is a graph isomorphism between them that preserves gains.

The *simplification* of a gain graph is obtained by removing all but one of each set of parallel edges, including parallel loops, that have the same gain. (Unlike the usual definition, we do not mean to remove all loops.)

A *circle* is a connected 2-regular subgraph, or its edge set. The gain of $C = e_1e_2 \cdots e_l$ is $\varphi(C) := \varphi(e_1)\varphi(e_2) \cdots \varphi(e_l)$; this is not entirely well defined, but it is well defined whether the gain is or is not the neutral element of \mathfrak{G} . An edge set or subgraph is called *balanced* if every circle in it has neutral gain.

Switching Φ by a *switching function* $\eta : V \rightarrow \mathfrak{G}$ means replacing φ by φ^η defined by

$$\varphi^\eta(e_{ij}) := \eta_i^{-1}\varphi(e_{ij})\eta_j,$$

where $\eta_i := \eta(v_i)$. We write Φ^η for the switched gain graph (Γ, φ^η) . Switching does not change balance of any subgraph.

The operation of deleting an edge or a set of edges is obvious, and so is contraction of a neutral edge set S : we identify each block $W \in \pi(S)$ to a single vertex and delete S while retaining the gains of the remaining edges. (This is just as with ordinary graphs.) The contraction is written Φ/S . A *neutral-edge minor* of Φ is any graph obtained by deleting and contracting neutral edges; in particular, Φ is a neutral-edge minor of itself.

Contraction of a general balanced edge set is not so obvious. Let S be such a set. There is a switching function η such that $\varphi^\eta|_S = 1_{\mathfrak{G}}$ [11, Section I.5]; and η is determined by one value in each component of S . If the endpoints of e are joined by a path P in S , $\varphi^\eta(e)$ is well defined as the value $\varphi(P \cup e)$, the gain of the circle formed by P and e . Now, to contract S we first switch so that S has all neutral gains; then we contract S as a neutral edge set. This contraction is also written Φ/S . It is well defined only up to switching because of the arbitrary choice of η . However, when contracting a neutral edge set we always choose $\eta \equiv 1_{\mathfrak{G}}$; then Φ/S is completely well defined.

An *edge minor* of Φ is a graph obtained by any sequence of deletions and contractions of any edges; thus, for instance, Φ is an edge minor of itself.

Contracting Φ by a *partition* π of V means identifying each block of π to a single vertex without changing the gain of any edge. The notation for this contraction is Φ/π . (A contraction of Φ by a partition is not an edge minor.)

If $W \subseteq V$, the subgraph induced by W is written $\Gamma:W$ or $\Phi:W$. For $S \subseteq E$, $S:W$ means the subset of S induced by W .

3 Chromatic relations on graphs

We begin with weak chromatic functions of ordinary graphs. This introduces the ideas in the relatively simple context of graphs; we also use the graph case in proofs. The models are the chromatic polynomial, $\chi_\Gamma(q)$, and its normalized derivative the beta invariant $\beta(\Gamma) := (-1)^n(d/dq)\chi_\Gamma(1)$. Regarded as a function F of graphs, a weak chromatic function has three fundamental properties: the *deletion-contraction* law,

$$F(\Gamma) = F(\Gamma \setminus e) - F(\Gamma/e) \quad \text{for every link } e;$$

loop nullity,

$$F(\Gamma) = 0 \quad \text{if } \Gamma \text{ has a loop};$$

and *invariance under simplification*,

$$F(\Gamma') = F(\Gamma) \quad \text{if } \Gamma' \text{ is obtained from } \Gamma \text{ by simplification.}$$

Two usually important properties of which we have no need are isomorphism invariance (with the obvious definition) and multiplicativity (the value of F equals the product of its values on components; the chromatic polynomial is multiplicative but the beta invariant is not).

Let \mathcal{G} be the class of all graphs. A function F from \mathcal{G} to an abelian group is a *weak chromatic function* if it satisfies deletion-contraction, invariance under simplification, and loop nullity. (Chromatic functions are a special type of Tutte function, as defined in [12]. The lack of multiplicativity is why our functions are “weak”; cf. [12].) One can take a function F with smaller domain, in particular, the set $\mathcal{M}(\Gamma_0)$ of all *edge minors* of a fixed graph Γ_0 (these are the graphs obtained from Γ_0 by deleting and contracting edges), although then Lemma 3.1 is less strong.

The first result shows that loop nullity is what particularly distinguishes chromatic functions from other functions that satisfy the deletion-contraction property, such as the number of spanning trees.

Lemma 3.1. *Given a function F from all graphs to an abelian group that satisfies deletion-contraction for all links, loop nullity is equivalent to invariance under simplification.*

Let Γ_0 be any graph. Given a function F from $\mathcal{M}(\Gamma_0)$ to an abelian group that satisfies deletion-contraction for all links, loop nullity implies invariance under simplification.

Proof. Assume F has loop nullity. Suppose we contract an edge in a digon $\{e, f\}$. The other edge becomes a loop. By deletion-contraction and loop nullity, $F(\Gamma) = F(\Gamma \setminus e) - F(\Gamma/e) = F(\Gamma \setminus e) + 0$.

Conversely, assume F is invariant under simplification. If Γ' has a loop f , it is the contraction of a graph Γ with a digon $\{e, f\}$ and the preceding reasoning works in reverse. This reasoning applies if the domain of F contains a graph Γ with the requisite digon; that is certainly true if the domain is \mathcal{G} . \square

In other words, one may omit invariance under simplification from the definition of a weak chromatic function. This is good to know because for some functions it is easier to establish loop nullity than invariance under simplification.

A weak chromatic function can be treated as a function of vertex-labelled graphs. (A *vertex-labelled (simple) graph* has a vertex set and a list of adjacent vertices.) To see this we need only observe that graphs with loops can be ignored, since they have value $F(\Gamma) = 0$ by definition, and for graphs without loops, the set of vertices and the list of adjacent pairs determines $F(\Gamma)$. To prove the latter, consider two simple graphs, Γ_1 and Γ_2 , with the same vertex set and adjacencies. Let Γ be their union (on the same vertex set). Both Γ_1 and Γ_2 are simplifications of Γ , so all three have the same value of F . When we regard F as defined on vertex-labelled graphs, contraction is *simplified contraction*; that is, parallel edges in a contraction Γ/e are automatically simplified.

Now we introduce the algebraic formalism of weak chromatic functions. The *chromatic group* for graphs, \mathbf{C} , is the free abelian group $\mathbb{Z}\mathcal{G}$ generated by all graphs, modulo the relations implied by deletion-contraction, invariance under simplification, and loop nullity. These relations are:

$$\begin{aligned} \Gamma &= (\Gamma \setminus e) - (\Gamma/e) && \text{for a link } e, \\ \Gamma' &= \Gamma && \text{if } \Gamma' \text{ is obtained from } \Gamma \text{ by simplification,} \\ \Gamma &= 0 && \text{if } \Gamma \text{ has a loop.} \end{aligned} \tag{3.1}$$

The point of these relations is that any homomorphism from the chromatic group to an abelian group will be a weak chromatic function of graphs, and any such function of graphs that has values in an abelian group \mathfrak{A} is the restriction to \mathcal{G} of a (unique) homomorphism from \mathbf{C} to \mathfrak{A} . These facts follow automatically from the definition of the chromatic group. Thus, \mathbf{C} is the universal abelian group for weak chromatic functions of graphs.

It follows from the definition (the proof is similar to the preceding one for functions) that two graphs are equal in \mathbf{C} if they simplify to the same vertex-labelled graph; that is, we may treat \mathbf{C} as if it were generated by vertex-labelled graphs and contraction as simplified contraction.

We may replace \mathcal{G} by $\mathcal{M}(\Gamma)$; $\mathbf{C}(\Gamma)$ denotes the corresponding chromatic group, i.e., $\mathbb{Z}\mathcal{M}(\Gamma)$ modulo the relations (3.1). Then $\mathbf{C}(\Gamma)$ is the universal abelian group for weak chromatic functions defined on the edge minors of Γ .

Now we present the main result about graphs. Let $\mathcal{F}(\Gamma_0)$ be the lattice of closed edge sets of the ordinary graph Γ_0 , and write μ for its Möbius function. If the empty set is not closed (that is, if there are loops in Γ_0), then $\mu(\emptyset, S)$ is defined to be identically 0.

Lemma 3.2. *In the chromatic group $\mathbf{C}(\Gamma_0)$ of a graph Γ_0 , for any edge minor Γ of Γ_0 we have the relations*

$$\Gamma = \sum_{S \in \mathcal{F}(\Gamma)} \mu(\emptyset, S) [(\Gamma/S) \setminus S^c] = \sum_{S \subseteq E} (-1)^{|S|} [(\Gamma/S) \setminus S^c]$$

and

$$\Gamma = \sum_{\pi \in \Pi^*(\Gamma)} K_\pi.$$

Proof. The two sums in the first identity are equal by Lemma 2.1 and because $(\Gamma/S)\setminus E = (\Gamma/R)\setminus E$.

We prove the first identity for graphs Γ without loops by induction on the number of edges in Γ . Let e be a link in Γ .

$$\begin{aligned}
\Gamma &= (\Gamma\setminus e) - (\Gamma/e) \\
&= \sum_{S\subseteq E\setminus e} (-1)^{|S|} [(\Gamma/S)\setminus S^c] - \sum_{A\subseteq E\setminus e} (-1)^{|A|} [(\Gamma/e)/A\setminus A^c] \\
&= \sum_{e\notin S\subseteq E} (-1)^{|S|} [(\Gamma/S)\setminus S^c] - \sum_{e\in S\subseteq E} (-1)^{|S|-1} [(\Gamma/S)\setminus S^c] \\
&= \sum_{S\subseteq E} (-1)^{|S|} [(\Gamma/S)\setminus S^c],
\end{aligned}$$

where in the middle step we replaced $A\subseteq E\setminus e$ by $S = A\cup e$.

If Γ has loops, let e be a loop. Then $\Gamma = 0$; at the same time

$$\begin{aligned}
\sum_{S\subseteq E} (-1)^{|S|} [(\Gamma/S)\setminus S^c] &= \sum_{e\in S\subseteq E} (-1)^{|S|} [(\Gamma/S)\setminus S^c] + \sum_{e\notin S\subseteq E} (-1)^{|S|} [(\Gamma/S)\setminus S^c] \\
&= \sum_{e\notin S\subseteq E} \left[(-1)^{|S|+1} [(\Gamma/S/e)\setminus (S\setminus e)^c] + (-1)^{|S|} [(\Gamma/S)\setminus S^c] \right],
\end{aligned}$$

which equals 0 because for a loop, $\Gamma/S/e = \Gamma/S\setminus e$.

We prove the second identity by induction on the number of edges not in Γ . Each stable partition of $\Gamma\setminus e$ is either a stable partition of Γ , or has a block that contains the endpoints of e . In that case, contracting e gives a stable partition of Γ/e . Thus,

$$\Gamma\setminus e = \Gamma + (\Gamma/e) = \sum_{\pi\in\Pi^*(\Gamma)} K_\pi + \sum_{\pi\in\Pi^*(\Gamma/e)} K_\pi = \sum_{\pi\in\Pi^*(\Gamma\setminus e)} K_\pi. \quad \square$$

Validity of the identities in $\mathbf{C}(\Gamma_0)$ implies they are valid in \mathbf{C} , since the former maps homomorphically into the latter by linear extension of the natural embedding $\mathcal{M}(\Gamma_0) \rightarrow \mathcal{G}$. (We cannot say that this homomorphism of chromatic groups is injective. The relations in \mathbf{C} might conceivably imply relations amongst the minors of Γ_0 that are not implied by the defining relations of $\mathbf{C}(\Gamma_0)$. We leave the question of injectivity open since it is not relevant to our work.)

4 Neutral chromatic functions and relations on gain graphs

We are interested in functions on gain graphs, with values in some fixed abelian group, that satisfy close analogs of the chromatic laws for graphs, which in view of Lemma 3.1

are two: deletion-contraction and loop nullity. The *neutral deletion-contraction* relation is the deletion-contraction identity

$$F(\Phi) = F(\Phi \setminus e) - F(\Phi / e) \tag{4.1}$$

for the special case where e is a neutral link. *Neutral-loop nullity* is the identity

$$F(\Phi) = 0 \text{ if } \Phi \text{ has a neutral loop.}$$

Neutral deletion-contraction is a limited version of a property that in the literature is usually required (if at all) of all or nearly all links. We call a function that adheres to these two properties a *weak neutral chromatic function* of gain graphs; “weak” because it need not be multiplicative, “neutral” because only neutral edges must obey the two laws. (To readers familiar with the half and loose edges of [11]: a loose edge is treated like a neutral loop.)

A function is *invariant under neutral-edge simplification* if

$$F(\Phi) = F(\Phi') \text{ when } \Phi' \text{ is obtained from } \Phi \text{ by removing one edge of a neutral digon.}$$

Lemma 4.1. *Given a function F of all gain graphs with a fixed gain group that satisfies deletion-contraction for all neutral links, neutral-loop nullity is equivalent to invariance under neutral-edge simplification.*

The proof is like that of Lemma 3.1 so we omit it.

The *neutral chromatic group* for \mathfrak{G} -gain graphs, $\mathbf{C}_0(\mathfrak{G})$, is the free abelian group $\mathbb{Z}\mathcal{G}(\mathfrak{G})$ generated by the class $\mathcal{G}(\mathfrak{G})$ of all gain graphs with the gain group \mathfrak{G} , modulo the relations implied by deletion-contraction of neutral links and neutral-loop nullity. These relations are

$$\begin{aligned} \Phi &= (\Phi \setminus e) - (\Phi / e) \text{ for a neutral link } e, \\ \Phi &= 0 \text{ if } \Phi \text{ has a neutral loop.} \end{aligned} \tag{4.2}$$

As with graphs, the purpose of these relations is that any homomorphism from the neutral chromatic group to an abelian group will be a function of \mathfrak{G} -gain graphs that satisfies neutral deletion-contraction and neutral-loop nullity, and every function of \mathfrak{G} -gain graphs that satisfies those two properties, and has values in an abelian group \mathfrak{A} , is the restriction of a (unique) homomorphism from $\mathbf{C}_0(\mathfrak{G})$ to \mathfrak{A} . (These facts follow automatically from the definition of the neutral chromatic group.) Thus, $\mathbf{C}_0(\mathfrak{G})$ is the universal abelian group for functions satisfying the two properties.

In the neutral chromatic group we get relations between gain graphs, in effect, by deleting and contracting neutral edges to expand any gain graph in terms of gain graphs with no neutral edge, while by addition and contraction we express it in terms of gain graphs whose neutral subgraph is the spanning complete graph $1_{\mathfrak{G}}K_n$.

Recall that $\mathcal{F}(\Gamma_0)$ is the lattice of closed sets of Γ_0 and $\Pi^*(\Gamma_0)$ is the set of stable partitions. Let μ_0 be the Möbius function of $\mathcal{F}(\Gamma_0)$. Recall also that, for a gain graph Φ , Γ_0 is the neutral subgraph of Φ and E_0 is the edge set of Γ_0 .

Theorem 4.2. *In the neutral chromatic group $\mathbf{C}_0(\mathfrak{G})$ we have the relation*

$$\Phi = \sum_{S \in \mathcal{F}(\Gamma_0)} \mu_0(0, S)[(\Phi/S) \setminus E_0] = \sum_{S \subseteq E_0} (-1)^{|S|} [(\Phi/S) \setminus E_0].$$

Theorem 4.3. *In the neutral chromatic group $\mathbf{C}_0(\mathfrak{G})$ we have the relation*

$$\Phi = \sum_{\pi \in \Pi^*(\Gamma_0)} [(\Phi/\pi) \cup 1_{\mathfrak{G}} K_\pi].$$

One can easily give direct proofs of Theorems 4.2 and 4.3 just like those of the two identities in Lemma 3.2. We omit these proofs in favor of ones that show the theorems are natural consequences of the relations for ordinary graphs; thus we deduce them by applying a homomorphism to the relations in Lemma 3.2.

Homomorphic Proof of Theorem 4.2. All vertices and edges are labelled, i.e., identified by distinct names. The vertices of a contraction are labelled in a particular way: the contraction by an edge set S has vertex set $\pi(S)$, the partition of V into the vertex sets of the connected components of (V, S) , and its edge set is the complement S^c of S . If we contract twice, say by (disjoint) subsets S and S' , then we label the vertices of the contraction as if $S \cup S'$ had been contracted in one step.

Now, define a function $f : \mathcal{M}(\Gamma_0) \rightarrow \mathcal{G}(\mathfrak{G})$ by

$$f(\Gamma_0/S \setminus T) := \Phi/S \setminus T.$$

Given an edge minor $\Gamma_0/S \setminus T$ of the neutral subgraph Γ_0 , even though we cannot reconstruct S and T separately, we can reconstruct the vertex partition $\pi(S)$ by looking at the labels of the vertices of the minor, and we can reconstruct $S \cup T$ by looking at the surviving edges of the minor. It follows that f is well defined, because its value on a minor of Γ_0 does not depend on which edges are contracted and which are deleted, as long as the vertex partition and surviving edge set are the same. (One can write this fact as a formula: $\Gamma_0/S \setminus T = \Gamma_0/\pi(S) \setminus (S \cup T)$.)

Extend f linearly to a function $\mathbb{Z}\mathcal{M}(\Gamma_0) \rightarrow \mathbb{Z}\mathcal{G}(\mathfrak{G})$ and define $\bar{f} : \mathbb{Z}\mathcal{M}(\Gamma_0) \rightarrow \mathbf{C}_0(\mathfrak{G})$ by composing with the canonical mapping $\mathbb{Z}\mathcal{G}(\mathfrak{G}) \rightarrow \mathbf{C}_0(\mathfrak{G})$. The kernel of \bar{f} contains all the expressions $G - [(G \setminus e) - (G/e)]$ for links e of edge minors $G \in \mathcal{M}(\Gamma_0)$ and all expressions G for edge minors with loops, because \bar{f} maps them all to $0 \in \mathbf{C}_0(\mathfrak{G})$ due to (4.2). Therefore, \bar{f} induces a homomorphism $F : \mathbf{C}(\Gamma_0) \rightarrow \mathbf{C}_0(\mathfrak{G})$.

Applying F to the first formula of Lemma 3.2, we get

$$F(\Gamma_0) = \sum_{S \in \mathcal{F}(\Gamma_0)} \mu_0(0, S) F((\Phi/S) \setminus E_0).$$

This is the theorem. □

Homomorphic Proof of Theorem 4.3. Apply the same F to the second formula of Lemma 3.2. □

If Φ is a gain graph, let $\Phi^0 := \Phi \cup 1_{\mathfrak{G}}K_n$, i.e., Φ with all possible neutral links added in.

Corollary 4.4. *In the neutral chromatic group, if Φ has no neutral edges, then*

$$\Phi^0 = \sum_{\pi \in \Pi_n} \mu(0, \pi)(\Phi/\pi),$$

where μ is the Möbius function of Π_n , and

$$\Phi = \sum_{\pi \in \Pi_n} (\Phi/\pi)^0.$$

Proof. Indeed, the identities follow from Theorems 4.2 and 4.3. In the theorems the graph of neutral edges of Φ^0 is the complete graph. So the flats are exactly the partitions of $[n]$. Contracting Φ by π introduces no neutral edges so in Theorem 4.2 it is not necessary to delete them. \square

5 Weak chromatic invariants of gain graphs

One can strengthen the definition of a weak chromatic function by requiring it to be invariant under some transformation of the gain graph. Examples:

- *Isomorphism invariance:* The value of F is the same for isomorphic gain graphs.
- *Switching invariance:* The value of F is not changed by switching.
- *Invariance under simplification:* F takes the same value on a gain graph and its simplification.

A *weak chromatic invariant* of gain graphs is a function that satisfies the deletion-contraction formula (4.1) for all links, neutral-loop nullity, and invariance under isomorphism, switching, and simplification. It is a *chromatic invariant* if it is also multiplicative on connected components, i.e.,

$$F(\Phi_1 \cup \Phi_2) = F(\Phi_1)F(\Phi_2). \tag{5.1}$$

Lemma 5.1. *Let F be a function on all gain graphs with a fixed gain group that is switching invariant and satisfies deletion-contraction for neutral links. Then neutral-loop nullity is equivalent to invariance under simplification and it implies isomorphism invariance.*

Proof. Suppose F is switching invariant and satisfies deletion-contraction for all neutral links. Then it satisfies deletion-contraction for all links, because any gain graph Φ can be switched to give gain $1_{\mathfrak{G}}$ to any desired link.

Suppose F also has neutral-loop nullity. When we contract an edge e in a balanced digon $\{e, f\}$, f becomes a neutral loop because of the switching that precedes contraction. By deletion-contraction and neutral-loop nullity,

$$F(\Phi) = F(\Phi \setminus e) - F(\Phi / e) = F(\Phi \setminus e) + 0.$$

Conversely, if Φ' has a neutral loop f , it is the contraction of a gain graph Φ with a neutral digon $\{e, f\}$ and the same reasoning works in reverse.

Now we deduce isomorphism invariance from invariance under simplification. Suppose we have two isomorphic gain graphs, Φ and Φ' , with different edge sets. We may assume the edge sets are disjoint and that, under the isomorphism, the vertex bijection is the identity and $e \leftrightarrow e'$ for edges. Take the union of the two graphs on the same vertex set; the union contains balanced digons $\{e, e'\}$. If we remove e from each pair we get Φ' , but if we remove e' we get Φ . The value of F is the same either way.

This treatment omits loops. Balanced loops make F equal to 0. For unbalanced loops we induct on their number. We can treat Φ as the contraction Ψ / f where Ψ is a gain graph in which e, f are parallel links and f is neutral, and similarly $\Phi' = \Psi' / f'$. Since $F(\Psi) = F(\Psi')$ and $F(\Psi \setminus e) = F(\Psi' \setminus e')$ by induction, $F(\Phi) = F(\Phi')$ by deletion-contraction. \square

Proposition 5.2. *A switching-invariant weak chromatic function of all gain graphs with a fixed gain group is a weak chromatic invariant.*

Proof. By switching we can make any link into a neutral link. Apply Lemma 5.1. \square

(We could formulate these properties of functions in terms of new chromatic groups, which are quotients of the weak chromatic group obtained by identifying gain graphs that are equivalent under a suitable equivalence, like simplification, switching, or isomorphism. However, that would contribute nothing to our general theory and it seems an overly complicated way to do the computations in the second half of this paper.)

Example 5.3. Weak chromatic invariants abound, but the most important is surely the total chromatic polynomial. A *multi-zero coloration* is a mapping $\kappa : V \rightarrow (\mathfrak{G} \times [k]) \cup [z]$, where k and z are nonnegative integers. It is *proper* if it satisfies none of the following edge constraints, for any edge e_{ij} :

$$\kappa(v_j) = \kappa(v_i) \in [z]$$

or

$$\kappa(v_i) = (m, g) \in \mathfrak{G} \times [k] \text{ and } \kappa(v_j) = (m, g\varphi(e_{ij})).$$

When \mathfrak{G} is finite, the *total chromatic polynomial* is defined by

$$\tilde{\chi}_\Phi(q, z) = \text{the number of proper multi-zero colorations,} \tag{5.2}$$

where $q = k|\mathfrak{G}| + z$. This function combines the *chromatic polynomial*,

$$\chi_{\Phi}(q) = \tilde{\chi}_{\Phi}(q, 1), \quad (5.3)$$

and the *zero-free chromatic polynomial*,

$$\chi_{\Phi}^*(q) = \tilde{\chi}_{\Phi}(q, 0), \quad (5.4)$$

of [10] and [11, Part III]. All three polynomials generalize the chromatic polynomial, for, regarding an ordinary graph Γ as a gain graph with gains in the trivial group, we see that

$$\tilde{\chi}_{\Gamma}(q, z) = \chi_{\Gamma}(q)$$

(which is independent of z), where $\chi_{\Gamma}(q)$ is the usual chromatic polynomial of Γ .

A second definition of the chromatic polynomials, which is algebraic, applies to all gain graphs, including those with infinite gain group. We define a total chromatic polynomial for any gain graph by the formula

$$\tilde{\chi}_{\Phi}(q, z) := \sum_{S \subseteq E} (-1)^{|S|} q^{b(S)} z^{c(S) - b(S)}, \quad (5.5)$$

where $b(S)$ is the number of components of (V, S) that are balanced, and we define the chromatic and zero-free chromatic polynomials by means of (5.3) and (5.4).

Proposition 5.4. *The total chromatic polynomial is a weak chromatic invariant of gain graphs. The combinatorial and algebraic definitions, (5.2) and (5.5), agree when both are defined.*

Proof. The first task is to show that the two definitions of the total chromatic polynomial agree. The combinatorial total chromatic polynomial is the special case of the state chromatic function $\chi_{\Phi}(\mathbf{Q})$ of [14, Section 2.2] in which the spin set $\mathbf{Q} = (\mathfrak{G} \times [k]) \cup [z]$. That is, $\tilde{\chi}_{\Phi}(k|\mathfrak{G}| + z, z) = \chi_{\Phi}(\mathbf{Q})$. This is obvious from comparing the definitions. Indeed, $\tilde{\chi}_{\Phi}(q, z)$ as defined in (5.2) is precisely the state chromatic function $\chi_{\Phi; \mathbf{Q}_1, \mathbf{Q}_2}(k_1, k_2)$ of the example in [14, Section 4.3] with the substitutions $q = k_1|\mathfrak{G}| + k_2$ and $z = k_2$.

Consequently, the combinatorial total chromatic polynomial has all the properties of a state chromatic function. The chief of these properties is that it agrees with the algebraic polynomial of (5.5). This fact is [14, Equation (4.3)] combined with Lemma 2.1 and the observation that the fundamental closure of S has the same numbers of components and of balanced components as does S [14, p. 144].

The second task is to prove that the algebraic total chromatic polynomial is a chromatic invariant. Isomorphism invariance is obvious from the defining equation (5.5). Switching invariance follows from the fact that $b(S)$ and $c(S)$ are unchanged by switching. Multiplicativity, Equation (5.1), is easy to prove by the standard method of splitting the sum over S into a double sum over $S \cap E(\Phi_1)$ and $S \cap E(\Phi_2)$. Reasoning like that in the proof of Lemma 3.2 proves neutral-loop nullity. By Proposition 5.2, if $\tilde{\chi}_{\Phi}$ satisfies

deletion-contraction, (4.1), for every link then it is invariant under simplification and thus is a chromatic invariant.

Thus, we must prove that $\tilde{\chi}_\Phi$ does satisfy deletion-contraction with respect to a link e . The method is standard—e.g., see the proof of [11, Theorem III.5.1]. We need two formulas about the contraction Φ/e . Suppose $e \in S \subseteq E$. Clearly, $c_\Phi(S) = c_{\Phi/e}(S \setminus e)$. [11, Lemma I.4.3] tells us that $b_\Phi(S) = b_{\Phi/e}(S \setminus e)$. Now we calculate:

$$\begin{aligned} \tilde{\chi}_\Phi(q, z) - \tilde{\chi}_{\Phi \setminus e}(q, z) &= \sum_{\substack{S \subseteq E \\ e \in S}} (-1)^{|S|} q^{b_\Phi(S)} z^{c_\Phi(S) - b_\Phi(S)} \\ &= \sum_{T \subseteq E \setminus e} (-1)^{|T|+1} q^{b_{\Phi/e}(T)} z^{c_{\Phi/e}(T) - b_{\Phi/e}(T)} \\ &= -\tilde{\chi}_{\Phi/e}(q, z), \end{aligned}$$

where again $T := S \setminus e$. This proves (4.1).

By Proposition 5.2, therefore, $\tilde{\chi}_\Phi$ is a chromatic invariant of gain graphs. \square

6 Integral gain graphs and integral affinographic hyperplanes

An *integral gain graph* is a gain graph whose gain group is the additive group of integers, \mathbb{Z} . The ordering of the gain group \mathbb{Z} singles out a particular switching function η_S : it is the one whose minimum value on each block of $\pi(S)$ is zero. We call this the *top switching function*. The contraction rule is that one uses the top switching function; thus the contraction can be uniquely defined, unlike the situation in general.

Contraction of a balanced edge set S in an integral gain graph can be defined quite explicitly.

First, we define η_S . In each component (V_i, S_i) of S , pick a vertex w_i and, for $v \in V_i$, define $\eta(v) := \varphi(S_{vw_i})$ for any path S_{vw_i} from v to w_i in S . (η is well defined because S is balanced.) Let v_i be a vertex which minimizes $\eta(v)$ in V_i . Define $\eta_S(v) := \varphi(S_{vv_i}) = \eta(v) - \eta(v_i)$ for $v \in V_i$. Then η_S is the top switching function for S , since $\eta_S(v_i) = 0 \leq \eta_S(v)$ for all $v \in V_i$.

Next, we switch. In Φ^{η_S} , the gain of an edge e_{vw} , where $v \in V_i$ and $w \in V_j$, is $\varphi^{\eta_S}(e_{vw}) = -\eta_S(v) + \varphi(e_{vw}) + \eta_S(w) = \varphi(S_{v_i v}) + \varphi(e_{vw}) + \varphi(S_{w v_j}) = \varphi(S_{v_i v} e_{vw} S_{w v_j})$. That is, $\varphi^{\eta_S}(e_{vw})$ is the gain of a path from v_i to v_j that lies entirely in S except for e_{vw} if that edge is not in S . (If e_{vw} is in S , its switched gain is 0, consistent with the fact that then $v_i = v_j$.)

Finally, we contract S . We can think of this as collapsing all of V_i into the single vertex v_i and deleting the edges of S , while not changing the gain of any edge outside S from its switched gain $\varphi^{\eta_S}(e_{vw}) = \varphi(S_{v_i v} e_{vw} S_{w v_j})$. (If there happens to be an edge $v_i v_j$, it will have the same gain in Φ/S as it did in Φ .)

A kind of invariance that will now become important is:

- *Loop independence*: The value of F is not changed by removing nonneutral loops.

Loop independence, when it holds true, permits calculations by means of contraction. The zero-free chromatic polynomial and both of the next two examples have loop independence, which we employ to good effect in Propositions 7.1 and 9.1.

Example 6.1. The *integral chromatic function* $\chi_{\Phi}^{\mathbb{Z}}(q)$ (from [3]) is the number of proper colorations of Φ by colors in the set $[q]$, *proper* meaning subject to the conditions given by the gains of the edges. This function is a weak chromatic function of integral gain graphs but it is not invariant under switching, so it is not a weak chromatic invariant. It is loop independent, because the color of a vertex is never constrained by a loop with nonzero additive gain.

That the integral chromatic function has neutral-loop nullity is obvious from the definition. To show it satisfies neutral deletion-contraction, consider a proper coloration of $\Phi \setminus e$, where e is a neutral link, using colors in $[q]$. If the endpoints of e have different colors, we have a proper coloration of Φ ; if they have the same color, we have a proper coloration of Φ/e . (This argument is standard in graph coloring, corresponding to the fact that the neutral subgraph acts like an ordinary graph.)

The reasoning fails if e has non-identity gain, and switching really does change $\chi_{\Phi}^{\mathbb{Z}}(q)$. Consider Φ with two vertices 1 and 2 and one edge e of nonnegative gain $g \in [q]$ in the orientation from 1 to 2. All such gain graphs are switching equivalent. The rule for a proper coloration κ is that $\kappa_2 \neq \kappa_1 + g$. Of all the q^2 colorations, the number excluded by this requirement is $q - g$ (or 0 if $q - g < 0$). Assuming $0 \leq g \leq q$, $\chi_{\Phi}^{\mathbb{Z}}(q) = q(q - 1) + g$, obviously not a switching invariant.

Example 6.2. The *modular chromatic function* $\chi_{\Phi}^{\text{mod}}(q)$ (also from [3]) is the number of proper colorations of the vertices by colors in \mathbb{Z}_q .

The remarks at the end of [3, Section 6] imply that $\chi_{\Phi}^{\text{mod}}(q)$ is a weak chromatic invariant. The idea is that $\chi_{\Phi}^{\text{mod}}(q) = \chi_{\Phi(\text{mod } q)}^*(q)$ where $\Phi(\text{mod } q)$ is Φ with gains modulo q . Take integral gain graphs Φ and Φ' such that Φ' is isomorphic to some switching of Φ . Then the same is true for $\Phi(\text{mod } q)$ and $\Phi'(\text{mod } q)$ with switching modulo q . Since $\chi_{\Phi(\text{mod } q)}^*(q)$ for fixed q is a weak chromatic invariant of gain graphs with gains in \mathbb{Z}_q , $\chi_{\Phi}^{\text{mod}}(q)$ is a weak chromatic invariant of integral gain graphs.

The modular chromatic function is loop independent, for the same reason as is the integral chromatic function.

The modular chromatic function is not too different from the zero-free chromatic polynomial. Write

$$\max_{\circ}(\Phi) := \text{the maximum gain of any circle in } \Phi.$$

Lemma 6.3 (see [13, Section 11.4, p. 339]). *The modular chromatic function of an integral gain graph Φ is given by*

$$\chi_{\Phi}^{\text{mod}}(q) = \chi_{\Phi}^*(q) \quad \text{for integers } q > \max_{\circ}(\Phi),$$

but equality fails in general for $q = \max_{\circ}(\Phi)$.

Proof. If we take the integral gains modulo $q > \max_{\circ}(\Phi)$, we do not change the balanced circles, because no nonzero circle gain is big enough to be reduced to 0. By Equation (5.5) with $z = 0$ so the sum may be restricted to balanced edge sets S , we do not change the zero-free chromatic polynomial. Proper colorations in \mathbb{Z}_q with modular gains are proper modular colorations.

If $q = \max_{\circ}(\Phi)$, at least one unbalanced circle becomes balanced so we do change the list of balanced edge sets and we cannot expect equality. (We have not tried to decide whether equality is possible at all.) \square

This lemma, though disguised by talk about finite fields and the Critical Theorem, is fundamentally the same method used by Athanasiadis in most of his examples in [2]; see [3, Section 6].

The *affinographic hyperplane arrangement* that corresponds to an integral gain graph Φ is the set \mathcal{A} of all hyperplanes in \mathbb{R}^n whose equations have the form $x_j = x_i + g$ for edges ge_{ij} in Φ . See [11, Section IV.4] or [3] for more detail about this connection. A most important point is that the characteristic polynomial of this arrangement, $p_{\mathcal{A}}(q)$, equals the zero-free chromatic polynomial $\chi_{\Phi}^*(q)$, by [11, Theorem III.5.2 and Corollary IV.4.5]. Examples include the well known Shi, Linial, and Catalan arrangements, which we will define. (In these definitions, \mathbb{Z} could be replaced by any ordered abelian group, or a subgroup of the additive group of any field; for instance, the additive real numbers.)

7 Catalan arrangements and their graphs

We will now apply the preceding results to obtain relations between the special gain graphs corresponding to the Shi, Linial, and Catalan arrangements. We begin with the last.

Let $C_n = \{0, \pm 1\}K_n$, the complete graph K_n (on vertex set $[n]$) with gains $-1, 0$, and 1 on every edge ij ; we call this the *Catalan graph* of order n , because the corresponding hyperplane arrangement is known as the Catalan arrangement, \mathcal{C}_n . Let $C'_n = \{\pm 1\}K_n$, the complete graph K_n with gains -1 and 1 on every edge ij ; we call this the *hollow Catalan graph*.

Let $c(n, j)$ be the number of permutations of $[n]$ with j cycles. The *Stirling number of the first kind* is $s(n, j) = (-1)^{n-j}c(n, j)$. The *Stirling number of the second kind*, $S(n, j)$, is the number of partitions of $[n]$ into j blocks.

Proposition 7.1. *Let F be a weak chromatic function of integral gain graphs with the property of loop independence. Between the Catalan and hollow Catalan graphs we have the two relations*

$$F(C'_n) = \sum_{j=1}^n S(n, j)F(C_j)$$

and

$$F(C_n) = \sum_{j=1}^n s(n, j)F(C'_j).$$

Proof. The proof begins with Corollary 4.4, which gives the identities

$$F(C_n) = \sum_{\pi \in \Pi_n} \mu(0, \pi) F(C'_n/\pi),$$

$$F(C'_n) = \sum_{\pi \in \Pi_n} F((C'_n/\pi)^0).$$

By the hypotheses on F , we can simplify the contractions. Contraction by a partition introduces no new gains, but only loops and multiple edges with the same gains. Multiple edges simplify without changing F because F is a weak chromatic function. The loops can be deleted because F is loop independent. Therefore, $F(C'_n/\pi) = F(C'_{|\pi|})$ and $F((C'_n/\pi)^0) = F(C_{|\pi|})$. It follows that

$$F(C_n) = \sum_{\pi \in \Pi_n} \mu(0, \pi) F(C'_{|\pi|}) = \sum_{j=1}^n s(n, j) F(C'_j)$$

because $s(n, j) = \sum_{\pi \in \Pi_n: |\pi|=j} \mu(0, \pi)$, and

$$F(C'_n) = \sum_{\pi \in \Pi_n} F(C_{|\pi|}) = \sum_{j=1}^n S(n, j) F(C_j). \quad \square$$

Let r_n be the number of regions in the Catalan arrangement \mathcal{C}_n and let r'_n be the number of regions of the arrangement corresponding to the hollow Catalan graph C'_j (which Stanley calls the *semiorder arrangement*).

Corollary 7.2. $r_n = \sum_j c(n, j) r'_j$.

Proof. The arrangement $\mathcal{H}[\Phi]$ corresponding to an additive real gain graph of order n has $(-1)^n \chi_{\Phi}^*(-1)$ regions [11, Corollary IV.4.5(b)]. Also, $(-1)^{n-j} s(n, j)$ is the number of i -cycle permutations. Applying the weak chromatic invariant χ^* to the second equation of Proposition 7.1 and evaluating at $q = -1$, we get $\chi_{C_n}^*(-1) = \sum_j s(n, j) \chi_{C'_j}^*(-1)$. Since $s(n, j) = (-1)^{n-j} c(n, j)$, the equation follows. \square

The integral, modular, and zero-free chromatic functions of C_n are very simple to obtain. The gains 0 correspond to the condition that the colors of the vertices are all different. The gains -1 and 1 correspond to the condition that the colors of two vertices are never consecutive. Thus, the modular chromatic function of C_n counts injections $f : V \rightarrow \mathbb{Z}_q$ such that no two values of f are consecutive. We repeat the well known evaluation [7]. If we shift f so that $f(v_1) = 0$ (thereby collapsing q different injections together) and delete the successor of each value, we have an injection $\bar{f} : V \rightarrow \mathbb{Z}_{q-n}$ such that $\bar{f}(v_1) = 0$, or equivalently, an arbitrary injection $\bar{f}' : V \setminus \{v_1\} \rightarrow \{2, \dots, q-n\}$. There are $(q-n-1)_{n-1}$ of these. It follows that

$$\chi_{C_n}^{\text{mod}}(q) = q(q-n-1)_{n-1} \quad \text{for integers } q > n$$

(and it obviously equals 0 for $q < 2n$). Similarly, the integral chromatic function is

$$\chi_{C_n}^{\mathbb{Z}}(q) = (q - n)_n \quad \text{for integers } q \geq n$$

(and 0 for $q < 2n$). The zero-free chromatic polynomial is

$$\chi_{C_n}^*(q) = q(q - n - 1)_{n-1}$$

by Lemma 6.3 (or see [13, Equation (11.1)], where $\chi_{C_n}^*(q)$ is called $\chi_{[-1,1]K_n}^b(q)$).

To obtain the various chromatic functions of the hollow Catalan graphs C'_n directly is not as easy, but they follow from Proposition 7.1. We observed in Example 6.1 that the integral chromatic function is a weak chromatic function and loop independent; therefore,

$$\chi_{C'_n}^{\mathbb{Z}}(q) = \sum_{j=1}^n S(n, j)(q - j)_j \quad \text{for } q \geq n.$$

Since the zero-free chromatic polynomial is a chromatic invariant with loop independence,

$$\chi_{C'_n}^*(q) = q \sum_{j=1}^n S(n, j)(q - j - 1)_{j-1}.$$

Then the modular chromatic function follows by Lemma 6.3 and the fact that $\max_{\odot}(C_n) = n$:

$$\chi_{C'_n}^{\text{mod}}(q) = q \sum_{j=1}^n S(n, j)(q - j - 1)_{j-1} \quad \text{for } q > n.$$

8 Arrangements between Shi and Catalan

The *Shi graph* of order n is $S_n = \{0, 1\}\vec{K}_n$, i.e., the complete graph K_n with gains 0 and 1 on every oriented edge ij with $i < j$. To have simple notation we take vertex set $[n]$ and we write all edges ij with the assumption that $i < j$.

Let G be a spanning subgraph of K_n , that is, it has all n vertices; and define $SC(G)$ to be the gain graph $S_n \cup \{-1\}\vec{G}$, which consists of the complete graph K_n with gains 0 and 1 on every edge ij , and also gain -1 if $ij \in E(G)$. We call $SC(G)$ a graph *between Shi and Catalan*. If G is edgeless we have the Shi graph S_n and if G is complete we have the Catalan graph C_n . We compute chromatic functions of these graphs between Shi and Catalan.

Let us start with the case where G has a unique edge $e_0 = i_0j_0$ and compute the modular chromatic function for large q (that is, the zero-free chromatic polynomial, by Lemma 6.3). The gains 0 correspond to the condition that the colors must be all different. The gains 1 on every edge correspond to the condition that if the color of i immediately

follows the color of j then $i < j$. Finally, the gain -1 on the edge e_0 implies that the color of i_0 cannot immediately follow the color of j_0 . The number of modular colorations of our graph can be deduced by taking out of the list of proper q -colorations of S_n the ones for which the colors of i_0 and j_0 are consecutive (in decreasing order). The number of these equals $q/(q-1)$ times the number of proper $(q-1)$ -colorations of S_{n-1} , since we can simply remove j_0 and its color; more precisely, we normalize the colorations so that i_0 has color 0, thus j_0 has color $q-1$, and then convert by removing j_0 and $q-1$; S_n becomes S_{n-1} and \mathbb{Z}_q becomes \mathbb{Z}_{q-1} . We conclude from the known value $\chi_{S_n}^*(q) = q(q-n)^{n-1}$ ([2, 6]; we reprove this soon) that the modular chromatic function of $SC(G)$ is $q[(q-n)^{n-1} - (q-n-1)^{n-1}]$.

This small example can make one feel the complexity of the computation in the case of a general graph G . Nevertheless one can produce formulas.

Let G be a graph on vertex set $[n]$. A *descending path* in G is a path $i_1 i_2 \cdots i_l$ such that $i_1 > i_2 > \cdots > i_l$. Let $p_r(G)$ be the number of ways to partition the vertex set $[n]$ into r blocks, each of which is the vertex set of a descending path. We call such a partition a *descending path partition* of G .

Proposition 8.1. *Let G be a simple graph with vertex set $[n]$. The integral and modular chromatic functions and the zero-free chromatic polynomial of $SC(G)$ are given by*

$$\begin{aligned}\chi_{SC(G)}^{\mathbb{Z}}(q) &= \sum_{r=1}^n p_r(G^c)(q-n+1)_r && \text{for } q \geq n-1, \\ \chi_{SC(G)}^{\text{mod}}(q) &= q \sum_{r=1}^n p_r(G^c)(q-n-1)_{r-1} && \text{for } q > n, \\ \chi_{SC(G)}^*(q) &= q \sum_{r=1}^n p_r(G^c)(q-n-1)_{r-1}.\end{aligned}$$

For the Shi graph in particular,

$$\begin{aligned}\chi_{S_n}^{\mathbb{Z}}(q) &= (q-n+1)^n && \text{for } q \geq n-1, \\ \chi_{S_n}^{\text{mod}}(q) &= q(q-n)^{n-1} && \text{for } q > n, \\ \chi_{S_n}^*(q) &= q(q-n)^{n-1}.\end{aligned}$$

As we mentioned, the formula for $\chi_{S_n}^*(q)$ is already known from Athanasiadis', Headley's, and Postnikov–Stanley's several computations of the characteristic polynomial of the Shi arrangement ([2, Theorem 3.3], [6], and [8, Example 9.10.1]), since the two polynomials are equal, as we remarked in Section 6.

Proof. We count proper colorations, extending the method Athanasiadis used to compute $\chi_{S_n}^{\text{mod}}(q)$ ([2], as reinterpreted in [3, Section 6]).

Our first remarks apply both to integral and modular coloring. The gains on the edges correspond to relations on the colors of the vertices. The 0-edges prevent coloring two different vertices with the same color. The 1-edges imply that if two vertices are colored

with consecutive colors then the larger vertex has the first color. This gives a nice way to describe permissible colorations in q colors.

Since all colors used are different we know that there are exactly $q - n$ colors not used (we must assume that $n \leq q$). Imagine these colors lined up in circular or linear order, depending on whether we are evaluating χ^{mod} or $\chi^{\mathbb{Z}}$. We now need to arrange the vertices in the spaces between these unused colors. When we place some vertices in the same space they are in descending order, so their places are compelled by their labels. Therefore, all we have to do for the Shi graph, where there are no -1 -edges, is to assign each of n vertices to a space between the $q - n$ unused vertices. This is the classical problem of placing labelled objects into labelled boxes. There are $q - n$ boxes in the modular case and $q - n + 1$ in the integral case. In the modular case, we assign vertex 1 to box 0; the other $n - 1$ vertices can be placed arbitrarily. Then after inserting the vertices we have a circular permutation of q objects, which is isomorphic to \mathbb{Z}_q in q ways; this accounts for the extra factor of q in the modular Shi formula.

For a gain graph between Shi and Catalan, the edges with gain -1 correspond to vertices $i < j$ that can be in the same box only if they are not consecutive amongst the vertices in the box; i.e., there must be present in the box at least one vertex h satisfying $i < h < j$. More precisely, suppose the vertices in the box, in descending order, are j_1, j_2, \dots, j_l . Then no consecutive pair can be adjacent in G , or, to put it differently, $j_1 j_2 \cdots j_l$ must be a path in G^c .

Now we count. We first look at modular coloring. Consider the colors not used to be null symbols labelled by \mathbb{Z}_{q-n} , i.e., these colors are cyclically ordered. To get the number of proper colorations, we choose a partition of $[n]$ into the vertex sets of r descending paths, and then we place the r paths, each one with its vertex set in descending order, into the $q - n$ spaces between the nulls. Due to the cyclic symmetry we can fix the space before $0 \in \mathbb{Z}_{q-n}$ to be the one where we put the path that contains vertex 1. There are $(q - n - 1)_{r-1}$ ways to place the other $r - 1$ paths. Now we have a cyclic arrangement of q objects, vertices and nulls. This set is isomorphic to \mathbb{Z}_q in q different ways, each of which gives a different proper coloration of $SC(G)$. We get for the modular chromatic polynomial

$$q \sum_{\mathcal{P}} (q - n - 1)_{r-1},$$

summed over all descending path partitions \mathcal{P} of G^c , where r is the number of paths in \mathcal{P} . Our description is meaningful so long as $q - n - 1 \geq 0$, since the largest possible number r is n .

For integral coloring the technique is similar. The nulls are linearly ordered, isomorphic to $[q - n]$, and there are $q - n + 1$ boxes, i.e., spaces between and around them. We get for the integral chromatic polynomial

$$\sum_{\mathcal{P}} (q - n + 1)_r$$

because the r paths can be placed in any distinct boxes. The computation applies as long as $q - n + 1 \geq 0$.

We derived the Shi formulas by a direct calculation but it is easy to deduce them from the general descending-path formulas. Since $G^c = K_n$, the descending-path-partition number $p_r(G^c)$ is just the number of partitions with r blocks, that is, $S(n, r)$. Then one can collapse the sums; e.g., for the zero-free chromatic polynomial,

$$\sum_{r=1}^n p_r(G^c)(q-n-1)_{r-1} = (q-n)^{-1} \sum_{r=1}^n S(n, r)(q-n)_r = (q-n)^{n-1}. \quad \square$$

When G^c is the comparability graph $\text{Comp}(P)$ of a partial ordering of $[n]$ that is compatible with the natural total ordering, i.e., such that $i <_P j$ implies $i < j$ (in \mathbb{Z}), a descending path is a chain, so $p_r(G^c)$ is the number of ways to partition the set P into r chains.

We will now limit ourselves to the special case where G is a graph of order k constructed from a partition of $[n]$. Let π partition $[n]$ into k blocks X_1, \dots, X_k , with the notation chosen so that $a_i := \min(X_1) < a_2 := \min(X_2) < \dots < a_k := \min(X_k)$. Thus, X_1 contains 1, X_2 contains the smallest element not in X_1 , and so on. The blocks are naturally partially ordered by letting $X_i < X_j$ if $c < d$ for every $c \in X_i$ and $d \in X_j$; in other words, if $b_i := \max X_i < a_j = \min X_j$. This partial ordering induces a partial order P_π on $[k]$. We say X_i and X_j *overlap* if neither $X_i < X_j$ nor $X_i > X_j$.

Let Γ_π be the interval graph of the intervals $[a_i, b_i]$ for $i \in [k]$. (See [5] for the many interesting properties of interval graphs.) Then Γ_π has an edge ij just when X_i overlaps X_j , so its complement is the comparability graph $\text{Comp}(P_\pi)$. The integral and modular chromatic functions of the gain graph $SC(\Gamma_\pi)$ for the partition π can be obtained directly in terms of Γ_π . Let $d_i = d_i(\pi)$ be the *lower degree* of i in Γ_π , i.e., the number of blocks X_j overlapping X_i and having $j < i$; note that $d_1 = 0$.

Theorem 8.2. *Let π be a partition of $[n]$ into k blocks X_1, \dots, X_k , and let $SC(\Gamma_\pi)$ be the corresponding gain graph (of order k) between Shi and Catalan. The integral and modular chromatic functions and the zero-free chromatic polynomial of $SC(\Gamma_\pi)$ are:*

$$\begin{aligned} \chi_{SC(\Gamma_\pi)}^{\mathbb{Z}}(q) &= (q-k+1) \prod_{i=2}^k (q-k+1-d_i) && \text{for } q \geq k-1 + \max d_i, \\ \chi_{SC(\Gamma_\pi)}^{\text{mod}}(q) &= q \prod_{i=2}^k (q-k-d_i) && \text{for } q \geq k + \max d_i, \\ \chi_{SC(\Gamma_\pi)}^*(q) &= q \prod_{i=2}^k (q-k-d_i). \end{aligned}$$

Direct Proof. Again we extend the method of Athanasiadis [2] used to compute $\chi_{S_n}^*(q)$ [2], but slightly differently from before.

We first look at modular coloring. To get the number of proper colorations, we choose the colors of the vertices in increasing order. To color the vertex X_1 , we have q choices. For X_2 we have $q-k$ choices if X_1 and X_2 do not overlap (which corresponds to the

presence of a -1 edge) and only $q - k - 1$ choices otherwise. We go on, and when coloring X_i we have $q - k - d_i$ choices (we use the fact that two blocks overlapping with X_i must overlap each other). We get for the modular chromatic polynomial

$$q \prod_{i=2}^k (q - k - d_i).$$

The lower bound on q arises from the fact that $q - k - d_i$ must never be negative if the reasoning is to hold good.

For integral coloring the technique is similar. To color X_1 we have $q - k + 1$ choices, which is the number of boxes. To color X_i we have $q - k + 1 - d_i$ choices, the number of boxes minus the number of forbidden boxes. We get for the integral chromatic polynomial

$$(q - k + 1) \prod_{i=2}^k (q - k + 1 - d_i) = \prod_{i=1}^k (q - k + 1 - d_i). \quad \square$$

Deduction from Proposition 8.1. A *simplicial vertex ordering* in a graph G is a numbering of the vertices by $1, 2, \dots, k$ such that, for each r , in the subgraph G_r induced by the vertices $1, \dots, r$ the neighborhood of r is a clique (see, for instance, [5]). In $G = \Gamma_\pi$ it is easy to see that the natural ordering of $[n]$ is a simplicial vertex ordering and the number of neighbors of r in G_r is the lower degree d_r . A descending path is a chain in P_π .

We apply Proposition 8.1 inductively, leaving the easy case $k = 1$ to the reader. Suppose it is true for $P_\pi \setminus k$. A chain decomposition of P_π is obtained by taking an r -chain decomposition of $P_\pi \setminus k$ and adjoining k in either of two ways: we can add a new chain $\{k\}$, or we can add k to an existing chain $i_1 > \dots > i_l$, necessarily at the top. This is possible if and only if $k > i_1$ in P_π . Since the non-neighbors of k form a clique, they are mutually incomparable. Thus, each one is in a separate chain. Each must be the top element of its chain because otherwise the top element would be $< k$ and the non-neighbor would also be $< k$ by transitivity. It follows that the number of chains to which k can be added is $r - d_k$. We conclude that

$$p_r(\Gamma_\pi^c) = p_{r-1}(\Gamma_\pi^c \setminus k) + (r - d_k)p_r(\Gamma_\pi^c \setminus k).$$

Now we compute the value of the right-hand side of an expression in Proposition 8.1; we do the integral chromatic function, the others being similar. From the lemma,

$$\begin{aligned} \chi_{SC(\Gamma_\pi)}^{\mathbb{Z}}(q) &= \sum_{r=1}^k p_r(\Gamma_\pi^c)(q - k + 1)_r \\ &= \sum_{r=2}^k p_{r-1}(\Gamma_\pi^c \setminus k)(q - k + 1)_{r-1}(q - k - [r - 2]) \\ &\quad + \sum_{r=1}^{k-1} (r - d_k)p_r(\Gamma_\pi^c \setminus k)(q - k + 1)_r \end{aligned}$$

because $p_0(\Gamma_\pi^c \setminus k) = p_k(\Gamma_\pi^c \setminus k) = 0$,

$$\begin{aligned}
&= (q - k + 1 - d_k) \sum_{r=1}^{k-1} p_r(\Gamma_\pi^c \setminus k)(q - k + 1)_r \\
&= (q - k + 1 - d_k) \chi_{SC(\Gamma_\pi^c \setminus k)}^{\mathbb{Z}}(q - 1) \\
&= (q - k + 1 - d_k) \prod_{i=1}^{k-1} ([q - 1] - [k - 1] + 1 - d_i)
\end{aligned}$$

by induction. This is the desired formula. \square

9 Linial arrangements

The *Linial graph* of order n is $L_n = 1\vec{K}_n$, i.e., the complete graph K_n with gains 1 on every oriented edge ij (that is, with $i < j$). We found the integral chromatic function of the Linial graph in [3], but here we have a new and different formula that also gives the modular and zero-free chromatic functions. A corollary is a new formula for the characteristic polynomial of the Linial hyperplane arrangement, different from (though not as simple as) that of Athanasiadis [2, Theorem 4.2] (differently proved in [8, Example 9.10.3]).

We expand the Linial gain graph in gain graphs between Shi and Catalan. The chromatic functions we are calculating are invariant under simplification that does not remove neutral loops and are not affected by non-neutral loops.

Corollary 9.1. *Let F be a weak chromatic function of integral gain graphs which is loop independent. Then*

$$F(L_n) = \sum_{\pi \in \Pi_n} F(SC(\Gamma_\pi)).$$

Proof. A straightforward application of Corollary 4.4 shows that

$$L_n = \sum_{S \in \mathcal{F}(\Gamma_0)} S_n/S.$$

The flat S corresponds to a partition π ; it is the union of neutral cliques on the blocks of π . The edges that remain after contraction are the neutral edges, which connect all the blocks of π forming a complete neutral subgraph of S_n/S , and the contractions of 1-edges. Number the blocks by least element and partially order as in the previous section. If $i < j$, then there is a 1-edge from X_i to X_j in the contraction. If furthermore $X_i < X_j$ in P_π , there is a (-1) -edge in the rising direction. We can simplify multiple edges with the same gain, by definition of F . We can ignore neutral loops because contracting S leaves none. Contraction makes a 1-loop at each contraction vertex $X \in \pi$ that has a 1-edge, but these do not affect the value of F . Therefore, $F(S_n/S) = F(SC(\Gamma_\pi))$. \square

We get the next result from Theorem 8.2 and Corollary 9.1.

Theorem 9.2. *The integral and modular chromatic functions of the Linial gain graph satisfy*

$$\chi_{L_n}^{\mathbb{Z}}(q) = \sum_{\pi \in \Pi_n} (q - k + 1) \prod_{i=2}^k (q - k + 1 - d_i) \quad \text{for } q \geq n - 1,$$

$$\chi_{L_n}^{\text{mod}}(q) = q \sum_{\pi \in \Pi_n} \prod_{i=2}^k (q - k - d_i) \quad \text{for } q \geq n,$$

$$\chi_{L_n}^*(q) = q \sum_{\pi \in \Pi_n} \prod_{i=2}^k (q - k - d_i),$$

respectively, where $k = |\pi|$ and d_i is the lower degree $d_i(\pi)$.

Proof. We get the lower bound on q in the modular polynomial from Lemma 6.3, since the circle with maximum gain is $12 \cdots n1$, using 1-edges in the upward direction and the 0-edge $n1$. The gain is $(n - 1)1 + 0$.

In the integral case the lower bound follows from the obvious necessity that $q - |\pi| + 1 \geq 0$ for every partition. \square

The zero-free chromatic polynomial, by a remark in Section 6, is also the characteristic polynomial of the Linial arrangement \mathcal{L}_n ; thus,

$$p_{\mathcal{L}_n}(q) = q \sum_{\pi \in \Pi_n} \prod_{i=2}^k (q - k - d_i).$$

This new formula contrasts with that of Athanasiadis:

$$p_{\mathcal{L}_n}(q) = \frac{q}{2} \sum_{j=0}^n \binom{n}{j} \left(\frac{q-j}{2} \right)^{n-1}.$$

Interesting enumerative conclusions might follow from the equality of these two expressions for the Linial polynomial, but that is too complicated to pursue here. At any rate, our formula has an interesting combinatorial aspect.

Example 9.3. Let $n = 6$ and $\pi = \{13, 25, 46\}$, so $k = 3$. Then Γ_π has edges 12 and 23 but not 13. We have $d_1 = 0, d_2 = d_3 = 1$. The formula for integral colorations is $(q - 2)(q - 3)^2$, which gives two proper colorations for $q = 4$. The colorations are given by the sequences 31o2 and 2o31, where the color corresponds to the position in the sequence, the numbers denote the vertices, and o denotes an unused color. Thus in the first sequence vertex 3 is colored 1, vertex 1 has color 2, color 3 is not used, etc. The coloration is built up by taking $q - n = 1$ unused color, forming the sequence o , which makes two boxes to place descending paths in Γ_{pi^c} , and placing the descending paths 31 and 2 in the first and second boxes, respectively.

We see that different partitions π can give the same terms because the term of a partition depends only on the lower degrees. We would like to know, for a given nondecreasing sequence $D = (d_1 = 0, d_2, \dots, d_k)$, the number $N_1(D)$ of partitions of $[n]$ which correspond to this sequence. Knowing these numbers will give us formulas like

$$\chi_{L_n}^{\text{mod}}(q) = q \sum_D N_1(D) \prod_{i=2}^k (q - k - d_i).$$

(In fact, although each term of this sum depends only on the elements of the sequence and not on their order, we would be pleased to know also the number $N_2(D)$ of naturally ordered partitions of $[n]$ which correspond to an arbitrary ordered sequence D .)

Define $D_1(\pi)$ to be the *increasing lower degree sequence* of Γ_π , which is the sequence of lower degrees written in non-decreasing order; and let $D_2(\pi)$ be the sequence of lower degrees in *vertex order*, i.e., where $d_i = d_i(\pi) =$ the lower degree of X_i . To recognize these sequences is not difficult. Call an *ascent* of a sequence $D = (d_1, d_2, \dots, d_k)$ any position $i \in [k-1]$ such that $d_{i+1} > d_i$. The *ascent set* is $A(D) := \{i \in [k-1] : d_{i+1} > d_i\}$.

Proposition 9.4. *A sequence $D = (d_1, \dots, d_k)$ is the vertex-order lower degree sequence $D_2(\pi)$ of the overlap graph Γ_π for some $\pi \in \Pi_n$ if and only if*

$$\begin{aligned} d_1 &= 0 \leq d_2, \dots, d_k, \\ d_{i+1} &\leq d_i + 1 \text{ for every } i \in [k-1], \text{ and} \\ n &\geq k + \text{the number of ascents of } D. \end{aligned}$$

Proof. In the proof let \sim denote adjacency in Γ_π .

First we prove the three conditions are necessary. Let π be any partition of $[n]$ into k blocks; let $D = D_2(\pi)$. It is clear that D has the first property. For the second, let $j < i \leq k$. If X_j overlaps X_{i+1} it also overlaps X_i ; thus, if $j \sim i+1$ then $j \sim i$. Consequently, $d_{i+1} \leq d_i + 1$ and equality holds only if $a_{i+1} < b_i$. The latter implies $b_i > a_i$. We conclude that, in D ,

$$d_i < d_{i+1} \implies |X_i| > 1. \tag{9.1}$$

Consequently, every ascent of D requires a block of size at least 2; this gives the lower bound on n .

For sufficiency of the three conditions we assume a sequence D has the three properties of the proposition. We construct a partition $\pi \in \Pi_{n_D}$ that realizes D , where $n_D := k + |A(D)|$. (For $n > n_D$ we simply add $n_D + 1, \dots, n$ to the block of π that contains n_D .)

Define $d_0, d_{k+1} := 0$ and put $[i, j] := \{i, i+1, \dots, j\}$ if $i \leq j$. Now, let $X_i = \{a_i, b_i\}$ where

$$\begin{aligned} m_i &:= \min\{j \geq i : d_{j+1} \leq d_i\}, \\ t_i &:= |A(D) \cap [i]|, \\ a_i &:= i + t_{i-1} - d_i, \end{aligned}$$

$$\begin{aligned} b_i &:= a_i + (m_i - i) + |A(D) \cap [i, m_i]| \\ &= m_i + t_{m_i} - d_i. \end{aligned}$$

Finally, let $\pi := \{X_1, \dots, X_k\}$. Regardless of whether π is a partition or not, it has an overlap graph Γ_π which is the interval graph of the intervals $[a_i, b_i]$, and if all the a_i are distinct there is a natural ordering of the vertices of Γ_π .

Lemma 9.5. *The class π constructed from D satisfies:*

- (i) π is a partition of $[n_D]$.
- (ii) X_k is the singleton block $\{a_k = n_D - d_k\}$.
- (iii) The upper neighborhood of i in Γ_π is $[i + 1, m_i]$.
- (iv) The lower neighbors of i in Γ_π are the vertices $i_\delta := \max\{j \in [i - 1] : d_j = \delta\}$ for $\delta = 0, 1, \dots, d_i - 1$.
- (v) π has vertex-order lower degree sequence $D_2(\pi) = D$.

Proof. Note that the t_i are weakly increasing. Part (ii) is obvious because $m_k = k$.

We show that $a_1 < a_2 < \dots < a_k$. This will imply that, if we linearly order the X_i by least element, the ordering is subscript order. We calculate the increment:

$$a_{i+1} - a_i = 1 + (t_i - t_{i-1}) + (d_i - d_{i+1}) = 1 + u_i + (d_i - d_{i+1}).$$

If i is not an ascent, then $d_i - d_{i+1} \geq 0$ and $u_i = 0$, so $a_{i+1} - a_i > 0$. If i is an ascent, then $d_i - d_{i+1} = -1$ and $u_i = 1$; again $a_{i+1} - a_i > 0$.

Now we prove (iii). The first step is an equivalence, in which we assume $i < j$:

$$i \sim j \iff a_j \leq b_i \iff j + t_{j-1} - d_j \leq m_i + t_{m_i} - d_i.$$

If $j \leq m_i$, then $d_j > d_i$ so $j + t_{j-1} - d_j < m_i + t_{m_i} - d_i$. If $j > m_i$, then $t_{j-1} \geq t_{m_i}$. Suppose $a_j \leq b_i$. Then $j - m_i \leq (j - m_i) + (t_{j-1} - t_{m_i}) \leq d_j - d_{m_i} \leq j - m_i$. Thus, the inequalities are all equalities. It follows that $d_j - d_{m_i} = j - m_i$, which makes m_i an ascent, but by definition m_i can never be an ascent. Note that in both cases, $a_j = b_i$ is impossible. Therefore, no b_i and a_j can be equal if $i < j$. Neither can b_j equal a_i , because $b_j \geq a_j > a_i$.

Recall that $n_D = k + t_k$. Let d_j be the last member of D that has value 0. Then $t_j = k$ so $b_j = n_D$. We prove that all other $b_i < b_j$. If $i < j$, then $t_i < j$ so $b_i < a_j \leq b_j$. If $i > j$, then d_i is positive; therefore, $b_i = m_i + t_{m_i} - d_i < m_i + t_{m_i} \leq k + t_k = b_j$.

Since n_D is the number of blocks plus the number of doubleton blocks, and we have shown that none of the a 's and b 's can be equal except when $a_i = b_i$, we conclude that π is a partition of $[n_D]$.

It remains to prove (iv), from which (v) follows at once. Let $j < i$. We know from (iii) that

$$j \sim i \iff m_j \geq i \iff d_{j+1}, d_{j+2}, \dots, d_i > d_j.$$

It is easy to see that the last property implies, and is implied by, the property that $j = i_\delta$ for some δ , specifically for $\delta = d_j$. \square

The lemma obviously implies Proposition 9.4. \square

Proposition 9.6. *A sequence $D = (d_1, \dots, d_k)$ is the increasing lower degree sequence $D_1(\pi)$ of Γ_π for some partition $\pi \in \Pi_n$ if and only if it satisfies*

$$\begin{aligned} 0 &= d_1 \leq d_2 \leq \dots \leq d_k, \\ d_{i+1} &\leq d_i + 1 \text{ for every } i \in [k-1], \\ n &\geq k + d_k. \end{aligned}$$

Proof. D is a lower degree sequence if and only if some rearrangement of it, D' , satisfies the conditions of Proposition 9.4. The first two properties follow from those of D' . Neglecting the exact value of n , we see from Proposition 9.4 that D is a lower degree sequence if and only if it is itself, without rearrangement, a vertex-order lower degree sequence. The lower bound on n follows from that fact, because the smallest possible number of ascents of any rearrangement of D is $d_k = \max d_i$. \square

10 The total chromatic polynomial

We did not try to evaluate the chromatic polynomial of the Catalan and other graphs, all the more the total chromatic polynomial, because the results do not appear to be very nice nor are they known to count anything interesting. However, the total chromatic polynomial can be computed in terms of the zero-free polynomial, extending the balanced expansion

$$\chi_\Phi(q) = \sum_{\substack{W \subseteq V \\ W^c \text{ stable}}} \chi_{\Phi;W}^*(q-1)$$

from [11, Theorem III.6.1]. Now we develop this computation to the extent that evaluating the polynomial for the Catalan, Shi, or Linial graph becomes a mechanical exercise.

Proposition 10.1. *Let Φ be a gain graph with underlying graph Γ . The total chromatic polynomial satisfies the expansion identity*

$$\tilde{\chi}_\Phi(q, z) = \sum_{W \subseteq V} \chi_{\Gamma;W^c}(z) \chi_{\Phi;W}^*(q-z).$$

Lemma 10.2. *If Φ_1, \dots, Φ_c are the components of Φ , then*

$$\tilde{\chi}_\Phi(q, z) = \tilde{\chi}_{\Phi_1}(q, z) \cdots \tilde{\chi}_{\Phi_c}(q, z).$$

Proof. This is obvious from the combinatorial definition. The proof from the algebraic definition is standard (and easy) and is therefore omitted. \square

Combinatorial Proof of the Proposition. Assume \mathfrak{G} is finite, k, z are nonnegative integers, and $q = k|\mathfrak{G}| + z$. We count the proper colorations of Φ by the color set $([k] \times \mathfrak{G}) \cup [z]$. We pick a vertex subset W to color by $[k] \times \mathfrak{G}$, leaving the complement to be colored by $[z]$. The number of ways to color W properly is $\chi_{\Phi;W}^*(q-z)$. The number of ways to color W^c properly is $\chi_{\Gamma;W^c}(z)$ because a coloration is proper if and only if no edge has the same $[z]$ -color at both ends. \square

Algebraic Proof. The algebraic proof applies to all gain graphs. If Φ has no links, then the components are the single vertices with their loops, if any, for which the expansion is obviously correct. Thus, we may apply induction on the number of links. Assume there is a link e .

We develop the right-hand side of the expansion identity. We split the sum according to three cases: $e \in E:W$, $e \in E:W^c$, and e not in either induced edge set. In the first case, $\chi_{\Phi:W}^*(q-z) = \chi_{(\Phi \setminus e):W}^*(q-z) - \chi_{(\Phi/e):W}^*(q-z)$. In the second case, $\chi_{\Gamma:W^c}(z) = \chi_{(\Gamma \setminus e):W^c}(z) - \chi_{(\Gamma/e):W^c}(z)$. Let us write this all out. The right-hand side equals

$$\begin{aligned} & \sum_{\substack{W \subseteq V \\ e \in E:W}} \chi_{\Gamma:W^c}(z) \chi_{(\Phi \setminus e):W}^*(q-z) + \sum_{\substack{W \subseteq V \\ e \in E:W^c}} \chi_{(\Gamma \setminus e):W^c}(z) \chi_{\Phi:W}^*(q-z) \\ & + \sum_{\substack{W \subseteq V \\ e \notin (E:W) \cup (E:W^c)}} \chi_{\Gamma:W^c}(z) \chi_{\Phi:W}^*(q-z) \\ & - \left[\sum_{\substack{W \subseteq V \\ e \in E:W}} \chi_{\Gamma:W^c}(z) \chi_{(\Phi/e):W}^*(q-z) + \sum_{\substack{W \subseteq V \\ e \in E:W^c}} \chi_{(\Gamma/e):W^c}(z) \chi_{\Phi:W}^*(q-z) \right] \end{aligned}$$

The first three terms are precisely those in the expansion of $\tilde{\chi}_{\Phi \setminus e}(q, z)$. As for the remaining two terms, let v_e be the vertex of Φ/e that results from contracting e . Then

$$\begin{aligned} & \sum_{\substack{W \subseteq V \\ e \in E:W}} \chi_{\Gamma:W^c}(z) \chi_{(\Phi/e):W}^*(q-z) + \sum_{\substack{W \subseteq V \\ e \in E:W^c}} \chi_{(\Gamma/e):W^c}(z) \chi_{\Phi:W}^*(q-z) \\ & = \sum_{\substack{X \subseteq V(\Phi/e) \\ v_e \in E(\Phi/e):X}} \chi_{\Gamma:X^c}(z) \chi_{(\Phi/e):X}^*(q-z) + \sum_{\substack{X \subseteq V(\Phi/e) \\ v_e \in E(\Phi/e):X^c}} \chi_{(\Gamma/e):X^c}(z) \chi_{\Phi:X}^*(q-z) \end{aligned}$$

because a set $X \subseteq V(\Phi/e)$ corresponds to a set $W = X \setminus \{v_e\} \cup \{v, w\}$ if v, w are the endpoints of e in Φ , so that e is in either $E:W$ or $E:W^c$. The last two sums are the expansion of $\tilde{\chi}_{\Phi/e}(q, z)$. \square

Call Φ *complete* if every two vertices are adjacent by at least one edge. Recall that n denotes the order of Φ .

Corollary 10.3. *If Φ is complete, let L be the set of vertices that support loops. Then the total chromatic polynomial expands with falling-factorial coefficients:*

$$\tilde{\chi}_{\Phi}(q, z) = \sum_{L \subseteq W \subseteq V} (z)_{n-|W|} \chi_{\Phi:W}^*(q-z).$$

Proof. The factor $\chi_{\Gamma:W^c}(z)$ equals zero if there is a loop in W^c and otherwise it is the chromatic polynomial of $K_{|W^c|}$. \square

Call Φ *uniform* if all induced subgraphs of the same order are isomorphic. Then we can write Φ_m for an induced subgraph of order m ; Φ_n is just Φ . A uniform gain graph that is not complete has no links at all. If it has any loop, it has the same number of loops with the same gains at every vertex.

Corollary 10.4. *Suppose Φ is connected and uniform and has no loops. Then*

$$\tilde{\chi}_\Phi(q, z) = n! \sum_{j=0}^n \binom{z}{n-j} \frac{1}{j!} \chi_{\Phi_j}^*(q-z).$$

In particular, for the chromatic polynomial we have

$$\chi_\Phi(q) = \chi_{\Phi_n}^*(q-1) + n \chi_{\Phi_{n-1}}^*(q-z).$$

If there is an infinite uniform sequence Φ_0, Φ_1, \dots we can write Corollary 10.4 in terms of the exponential generating function, $\text{XGF}(a_n) := \sum_{n=1}^{\infty} a_n t^n / n!$, as

$$\text{XGF}(\tilde{\chi}_{\Phi_n}(q, z)) = (1+t)^z \text{XGF}(\chi_{\Phi_n}^*(q-z)). \quad (10.1)$$

Since the Catalan, Shi, and Linial graphs satisfy the hypotheses of Corollary 10.4 and the graphs between Catalan and Shi are complete, it is a routine matter to evaluate their chromatic and total chromatic polynomials by means of Corollary 10.4 or Equation (10.1). The intermediate graphs $SC(G)$ fall under Corollary 10.4.

11 Final remarks: Extended arrangements; weighted gain graphs; biased graphs

There are various ways to extend the Linial and Shi arrangements to more hyperplanes; each one corresponds to an integral gain graph of the form $\{-l+1, -l+2, \dots, m\} \vec{K}_n$, where $l, m \geq 0$ and $m > -l$. For the characteristic polynomial the cases $l = 0, 1$ were treated by Athanasiadis [2, Section 4], the general case by Postnikov and Stanley [8, Section 9]. It would certainly be interesting to apply our results to these arrangements.

A *weighted integral gain graph* (Φ, h) is an integral gain graph Φ together with a function $h : V \rightarrow \mathbb{Z}$, called the *weight* function. By the contraction rule for weights, when contracting a neutral edge set S the contracted weights are $h_S(W) := \max_{v \in W} h(v)$, where W is the vertex set of a component of S . For a weighted integral gain graph (Φ, h) one can define an integral chromatic function $\chi_{(\Phi, h)}^{\mathbb{Z}}(q)$ similar to that of Φ ; indeed, our $\chi_{\Phi}^{\mathbb{Z}}(q)$ is $\chi_{(\Phi, 0)}^{\mathbb{Z}}(q)$ (all weights equal 0). If e is a neutral link, then (from [3], reinterpreted from the equivalent rooted integral gain graphs of that paper)

$$\chi_{(\Phi, h)}^{\mathbb{Z}}(q) = \chi_{(\Phi \setminus e, h)}^{\mathbb{Z}}(q) - \chi_{(\Phi/e, h_e)}^{\mathbb{Z}}(q).$$

This shows there are interesting functions F on other structures than ordinary graphs and gain graphs—in fact, on arbitrary weighted gain graphs as defined in [4]—to which the general reductions Theorems 4.2 and 4.3 can apply.

Another kind of generalization is to biased graphs [11]. The balanced circles in a gain graph form a *linear subclass*, which means that if a theta subgraph contains two balanced circles, then its third circle is balanced. A *biased graph* is a pair $\Omega = (\Gamma, \mathcal{B})$ where \mathcal{B} is any linear subclass of the circles of Γ . An edge set is called balanced when every circle belongs to \mathcal{B} . Thus, a biased graph is a combinatorial generalization of a gain graph. There are no distinguished neutral edges in a biased graph, but one can choose any balanced edge set to play the same role. Thus, a deletion-contraction formula like that of Theorem 4.2 exists. There are also addition-contraction results where one adds edges to a balanced subgraph, though since it is probably not true that one can extend any balanced subset to a balanced K_n (this is an open question), there would be no complete generalization of Theorem 4.3.

Acknowledgement

We thank the referee for astute refereeing that helped to perfect the exposition.

References

- [1] Martin Aigner, *Combinatorial Theory*. Springer, Berlin, 1979. MR 80h:05002. Zbl 415.05001.
- [2] Christos A. Athanasiadis, Characteristic polynomials of subspace arrangements and finite fields. *Adv. Math.* **122** (1996), 193–233. MR 97k:52012. Zbl 872.52006.
- [3] David Forge and Thomas Zaslavsky, Lattice point counts for the Shi arrangement and other affinographic hyperplane arrangements. *J. Combin. Theory Ser. A* **114** (1) (2007), 97–109. MR 2007i:52026. Zbl 1105.52014.
- [4] David Forge and Thomas Zaslavsky, Colorations, orthotopes, and a huge polynomial Tutte invariant of weighted gain graphs. Submitted.
- [5] Martin Charles Golumbic, *Algorithmic Graph Theory and Perfect Graphs*. Academic Press, New York, 1980. MR 81e:68081. Zbl 541.05054.
- [6] Patrick Headley, On a family of hyperplane arrangements related to the affine Weyl groups. *J. Algebraic Combin.* **6** (1997), 331–338. MR 98e:52010. Zbl 911.52009.
- [7] Irving Kaplansky, Solution of the ‘Problème des ménages’. *Bull. Amer. Math. Soc.* **49** (1943), 784–785. MR 5, 86l. Zbl 60, 29d (e: 060.02904).
- [8] Alexander Postnikov and Richard P. Stanley, Deformations of Coxeter hyperplane arrangements. *J. Combin. Theory Ser. A* **91** (2000), 544–597. MR 2002g:52032. Zbl 962.05004.

- [9] Gian-Carlo Rota, On the foundations of combinatorial theory: I. Theory of Möbius functions. *Z. Wahrsch. verw. Gebiete* **2** (1964), 340–368. MR 30 #4688. Zbl 121, 24f (e: 121.02406).
- [10] Thomas Zaslavsky, Signed graph coloring. *Discrete Math.* **39** (1982), 215–228. MR 84h:05050a. Zbl 487.05027.
- [11] Thomas Zaslavsky, 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. Zbl 714.05057; 763.05096; 857.05088; 1031.05034.
- [12] Thomas Zaslavsky, Strong Tutte functions of matroids and graphs. *Trans. Amer. Math. Soc.* **334** (1992), 317–347. MR 93a:05047. Zbl 781.05012.
- [13] Thomas Zaslavsky, Perpendicular dissections of space. *Discrete Comput. Geom.* **27** (2002), 303–351. MR 2003i:52026. Zbl 1001.52011.
- [14] Thomas Zaslavsky, Totally frustrated states in the chromatic theory of gain graphs. *Europ. J. Combin.* **30** (2009), 133–156.