The Coloring Ideal and Coloring Complex of a Graph

Einar Steingrímsson
Matematik CTH & GU
412 96 Göteborg
SWEDEN
einar@math.chalmers.se

Abstract

Let G be a simple graph on d vertices. We define a monomial ideal K in the Stanley-Reisner ring A of the order complex of the Boolean algebra on d atoms. The monomials in K are in one-to-one correspondence with the proper colorings of G. In particular, the Hilbert polynomial of K equals the chromatic polynomial of G.

The ideal K is generated by square-free monomials, so A/K is the Stanley-Reisner ring of a simplicial complex C. The h-vector of C is a certain transformation of the *tail* $T(n) = n^d - \chi(n)$ of the chromatic polynomial χ of G. The combinatorial structure of the complex C is described explicitly and it is shown that the Euler characteristic of C equals the number of acyclic orientations of G.

1 Introduction and preliminaries

Let G be a simple graph on d vertices. In this paper we construct a monomial ideal K in the face ring (Stanley-Reisner ring) A of the order complex of the Boolean algebra on d atoms, which is isomorphic to a cone over the first barycentric subdivision of a $(d-1)$-simplex. The monomials in K are in one-to-one correspondence with the proper colorings of G. The quotient of A by K is the face ring of a simplicial complex whose structure can be described explicitly.

The construction of the ideal K is based on a definition of Chung and Graham [6], whose purpose was to give a combinatorial interpretation to the coefficients of the chromatic polynomial $\chi(n)$ of G when written in the basis $\{\binom{n+k}{d}\}_{k=0,\ldots,d}$. It was shown by Chow [5] that this result can also be derived from a theorem of Stanley's concerning his *chromatic symmetric function* [10]. However, our construction does not seem to be much related to Stanley's chromatic symmetric function. In fact, our complexes are isomorphic for the two non-isomorphic graphs on $n \geq 4$ vertices and two edges, whereas Stanley's functions for these graphs are not equal. On the other hand, our complex *does* distinguish the two non-isomorphic (but chromatically equivalent) graphs on five vertices that Stanley's function does not distinguish.

As it happens, invariants related to colorings of a graph G often have connections to the acyclic orientations of G. In our case, we show that the Euler characteristic of the coloring complex equals the number of acyclic orientations of G.

It should also be mentioned that Brenti asked [3] whether there exists, for an arbitrary graph G, a standard graded algebra whose Hilbert polynomial equals the chromatic polynomial of G. This was answered in the affirmative by Almkvist [1], but his proof is non-constructive. The structure of such a graded algebra, however, will not necessarily be closely related to the colorings of G, since its monomials of degree less than the chromatic number of G cannot correspond to colorings of G.

Throughout this paper, unless otherwise specified, G is a graph on d vertices labeled by the elements of $[d] = \{1, 2, \ldots, d\}$, with no loops and no multiple edges. Frequently, but not always, we suppress G from the notation for simplicity.

A *path of length* k in G is a sequence $v_0, v_1, v_2, \ldots, v_k$ of vertices of G such that there is an edge between v_{i-1} and v_i, for each $i \in [k]$.

A *stable* set in G is a set of vertices with no edge between any pair.

Let V be the set of vertices of G. A *coloring of* G is a map $\phi : V \longrightarrow \mathbb{N}$ with $\phi(x) \neq \phi(y)$ if x and y are adjacent in G, that is, if there is an edge between x and y. Thus, we treat the natural numbers as colors and when referring to the ordering of colors, we mean the usual ordering on $\mathbb{N}$. (Observe that we omit the word "proper" from the definition of coloring, since we only consider proper colorings).

A *coloring of* G *with* n *colors*, or n*-coloring*, is a map $\phi : V \longrightarrow [n]$ with $\phi(x) \neq \phi(y)$ if x and y are adjacent. (Observe that ϕ need not be surjective.)

Definition 1 Let $S_1, S_2, \ldots, S_m$ be an ordered partition of the vertices of G. For $v \in G$ let $\ell(v)$ be the length of the longest path $v_{i_1}, v_{i_2}, \ldots, v_{i_p} = v$ (ending in v) in G such that $v_{i_j} \in S_{i_j}$ for each j and $i_1 < i_2 < \cdots < i_p$.

If $\pi = a_1 a_2 \cdots a_d$ is a permutation in the symmetric group $\mathcal{S}_d$ we let π induce the ordered partition of the vertices in G obtained by letting a_i constitute the i-th block (singleton) in the ordered partition. In accordance with the definition of $\ell(v)$ subject to an ordered partition, we then let $\ell(k)$ be the length of the longest path $a_{i_1}, a_{i_2}, \ldots, a_k$ (ending in a_k) in G such that $i_1 < i_2 < \cdots < k$.

The following definition is a variation of a definition of Chung and Graham in [6, §5].

Definition 2 The integer $k \in [0, \ldots, d-1]$ is a *cut in* π *(with respect to* G*)* if

1. $k = 0$, or
2. $\ell(k) < \ell(k+1)$, or
3. $\ell(k) = \ell(k+1)$ and $a_k < a_{k+1}$.

Definition 3 Let $\pi = a_1 a_2 \cdots a_d$ be a permutation with cuts $i_1 = 0, i_2, \ldots, i_k$. The G*-sequence* of π is the sequence of sets $S_1, S_2, \ldots, S_k$ where $S_j = \{a_{i_j+1}, a_{i_j+2}, \ldots, a_{i_{j+1}}\}$, for $j < k$, and $S_k = \{a_{i_k+1}, a_{i_k+2}, \ldots, a_d\}$. The *short* G*-sequence of* π is $S_1, S_2, \ldots, S_{k-1}$.

As an example, let G be the graph

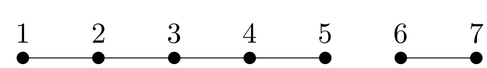

and let π be the permutation 5236417. Then the path lengths $\ell(k)$ associated to π are given by

$$\begin{array}{ccccccc} 5 & 2 & 3 & 6 & 4 & 1 & 7 \\ 0 & 0 & 1 & 0 & 2 & 1 & 1 \end{array}$$

so π has cuts 0, 2, 4 and 6 and thus G-sequence $\{2,5\},\ \{3,6\},\ \{1,4\},\ \{7\}$.

Lemma 4 *Let $S_1, S_2, \ldots, S_k$ be the G-sequence of a permutation π. If $a_m, a_{m+1} \in S_i$ then either $\ell(m) > \ell(m+1)$ or else $\ell(m) = \ell(m+1)$ and $a_m > a_{m+1}$. Moreover, each S_i is a stable set in G.*

Proof: The first part of the lemma follows directly from Definitions 2 and 3. Thus, if a_m and a_k both belong to S_i, with $m < k$, then $\ell(m) \geq \ell(k)$. But if a_m and a_k are adjacent in G then $\ell(m) < \ell(k)$, since a path ending in a_m can always be extended to end in a_k. That is a contradiction, so no two elements in S_i are adjacent in G. $\square$

The following theorem is stated (in a different, but equivalent, form) in [6], where it is claimed that it follows from the work of Brenti in [2]. Also, in [4] it is shown that the theorem follows from certain properties of the chromatic symmetric function of Stanley [10]. We give a proof that is different from both of these, but which better suits our purposes. First, a definition.

Definition 5 Let $P(n)$ be a polynomial of degree d. The *W-transform of P* is the polynomial W defined by

$$\sum_{n\geq 0} P(n)t^n = \frac{W(t)}{(1-t)^{d+1}}.$$

That W is a polynomial is easily shown, as is the fact that its degree is at most equal to the degree of P.

Theorem 6 *Let $W_G(t)$ be the W-transform of χ_G, that is,*

$$\sum_{n\geq 0} \chi_G(n)t^n = \frac{W_G(t)}{(1-t)^{d+1}}.$$

Then we have

$$W_G(t) = \sum_{\pi\in\mathcal{S}_d} t^{c(\pi)},$$

where $c(\pi)$ is the number of cuts in π. In particular, $\sum_{\pi\in\mathcal{S}_d} t^{c(\pi)}$ is independent of the labeling of G.

Proof: Define w_i, for $0 \le i \le d$, by setting $W_G(t) = w_0 + w_1 t + \cdots + w_d t^d$. Then the claim is equivalent to saying that

$$\chi_G(n) = \sum_{k=0}^{d} \binom{n+k}{d} w_{d-k},$$

or, in other words, that there are, for each permutation in $\mathcal{S}_d$ with $(d-k)$ cuts, exactly $\binom{n+k}{d}$ ways to color G with n colors. Let π be such a permutation. By the latter part of Lemma 4, there is an obvious way to associate a permutation π with k cuts to a coloring using k (out of n available) colors. Namely, color the vertices corresponding to the letters between the $(i-1)$-st and the i-th cut with color number i, and the letters after the last cut with color k.

Since π has $(d-k)$ cuts, there are exactly k places between adjacent letters in π that do not correspond to cuts. Let us pick, among n colors and k non-cuts, exactly d items, say i non-cuts and $(d-i)$ colors. Then we have chosen which $(d-i)$ colors to use, and which i non-cuts to retain. We now introduce extra cuts at the remaining $(k-i)$ non-cuts, which means we have a total of $(d-k)+(k-i) = (d-i)$ cuts, and thereby $(d-i)$ stable sets, which get colored by the $(d-i)$ colors chosen, in the order prescribed by π.

Thus, we have associated $\binom{n+k}{d}$ colorings to each permutation $\pi \in \mathcal{S}_d$ with $(d-k)$ cuts.

For the converse, we need to show that each coloring of G using some of n available colors arises from a unique permutation, together with a choice of extra cuts, as described above. Given such a coloring, partition the vertices of G into blocks, each consisting of all vertices with like color, and order the blocks increasingly by color. Within each block, order the vertices so that their corresponding maximal path lengths $\ell(m)$ are (weakly) decreasing, and so that the vertices are decreasing when two vertices have the same maximal path length associated to them. This can always be done, because the path lengths of two vertices in the same block depend only on the vertices in preceding blocks. Thus, writing the vertices in the order described we get a unique permutation π in $\mathcal{S}_d$.

By the construction of π, all of its cuts occur between blocks of the ordered partition P from which π was constructed. Thus, the coloring from which P was constructed arises from π together with the extra cuts (separating blocks in P) that are not cuts in π. $\square$

2 The coloring ideal

The field k in the following definition can be taken to be the complex numbers. Also, all rings are taken to be commutative. Recall that the short G-sequence of a permutation π is the G-sequence of π take away the last set in the sequence.

For undefined terminology and background in what follows, see [9].

Let $A = k\,[x_S \mid S \subseteq [d]]$, that is, A is the polynomial ring whose indeterminates correspond to all subsets of $[d]$. Throughout, R will denote the *face ring* (or *Stanley-Reisner ring*) of the order complex of the Boolean algebra on d atoms. This ring is the quotient A/I, where $I = \{x_S x_T \mid S \nsubseteq T \text{ and } T \nsubseteq S\}$. Thus, the monomials of R correspond precisely to those monomials $M = x_{S_1} x_{S_2} \cdots x_{S_k} \in A$ for which $\emptyset \subseteq S_1 \subseteq S_2 \subseteq \cdots \subseteq S_k \subseteq [d]$ (for some (unique) rearrangement of the indices).

Definition 7 Let $M = x_{S_1} x_{S_2} \cdots x_{S_k}$ be a monomial in A such that

$$\emptyset \subsetneq S_1 \subsetneq S_2 \subsetneq \cdots \subsetneq S_k \subsetneq [d],$$

where $\subsetneq$ denotes strict inclusion. Then M is a *basic coloring monomial for* G if there is a permutation $\pi \in \mathcal{S}_d$ with short G-sequence $S_1, S_2 \setminus S_1, \ldots, S_k \setminus S_{k-1}$.

A nonzero monomial $M = x_{S_1} x_{S_2} \cdots x_{S_k} \in A$ is a *coloring monomial for* G if M is divisible by a basic coloring monomial for G.

Definition 8 The *coloring ideal of* G is the ideal K_G in R generated by all (equivalently by the basic) coloring monomials for G.

Note that a coloring monomial may not be divisible by $x_{[d]}$, since no basic coloring monomial is divisible by $x_{[d]}$, as it is constructed from a *short* G-sequence. That we choose to define the coloring ideal in this way is due to a technicality which will be explained in Remark 15.

We shall now show that the monomials of degree n in K_G are in one-to-one correspondence with the colorings of G with $n+1$ colors.

It follows from the definitions of R and K_G that any monomial M in K_G can be written as

$$x_{S_1}^{e_1} x_{S_2}^{e_2} \cdots x_{S_k}^{e_k}$$

such that $S_1 \subsetneq S_2 \subsetneq \cdots \subsetneq S_k \subseteq [d]$ and such that $S_i \setminus S_{i-1}$ is a stable set in G for each i. Such a monomial gives rise to a unique coloring of G with n colors where n is the degree of M. Namely, if $S_1 = \emptyset$, the colors $1, 2, \ldots, e_1$ are not used. Otherwise, the vertices in S_1 get color 1. The vertices in $S_2 \setminus S_1$ get color $e_1 + 1$ and, in general, the vertices in $S_i \setminus S_{i-1}$ get color $e_1 + e_2 + \cdots + e_{i-1} + 1$. The vertices in $[d] \setminus S_k$ get color $\sum_i e_i + 1$. If $S_k = [d]$ and $e_k > 0$ then the last e_k colors are not used (recall that in an n-coloring not all n colors have to be used).

As an example, suppose

$$M = x_{\emptyset}^2 \cdot x_{25}^3 \cdot x_{235}^2$$

is a coloring monomial for G (where we write 25 for the set $\{2, 5\}$ etc.). In the corresponding coloring of G with 8 colors, the vertices 2 and 5 get color 3, the vertex 3 gets color 6 and all the remaining vertices (however many they are) get color 8. Multiplying M by $x_{[d]}^e$ corresponds to regarding the coloring in question as a coloring with $8+e$ colors. Clearly, two different monomials yield different colorings. If they are identical except for the exponent to $x_{[d]}$ they correspond to colorings with a different number of colors, although each vertex gets the same color in each of the two colorings.

Conversely, suppose we have a coloring of G with n colors (of which not all have to be used). Then each color is associated to a stable set. Order these sets increasingly by color and let S_i be the union of the first i of them. Then we can construct a corresponding coloring monomial as we now explain by an example. Suppose we have a 9-coloring of G where vertices 3 and 6 have color 4, vertex 7 has color 6 and vertices 1, 2, 4 and 5 have color 7. Then we have the sequence of sets $\{3,6\}$, $\{3,6,7\}$ and $\{1,2,3,4,5,6,7\}$. Given the colors used, and the fact that this is to be a 9-coloring, the corresponding coloring monomial is

$$x_\emptyset^3 \cdot x_{36}^2 \cdot x_{367} \cdot x_{[7]}^2.$$

We record this in the following theorem.

Theorem 9 *There is a one-to-one correspondence between the $(n+1)$-colorings of G and the monomials of degree n in K_G.* $\square$

As an immediate consequence we now get the following.

Corollary 10 *Let $F(K_G,t)$ be the Hilbert series of K_G and let d be the number of vertices in G. Then*

$$F(K_G,t) = \frac{1}{t} \cdot \frac{W_G(t)}{(1-t)^{d+1}},$$

Equivalently, $H(K_G,n) = \chi_G(n+1)$, that is, the Hilbert polynomial of K_G equals, up to a shift by one, the chromatic polynomial of G. $\square$

3 The coloring complex and its face ring

Clearly, the basic coloring monomials are square-free. Thus, the quotient R/K_G is the face ring of a simplicial complex whose vertex set is a subset of $\{S \mid S \subseteq [d]\}$ and whose minimal non-faces are

$$\{\ \{S_1, S_2, \ldots, S_k\} \mid x_{S_1} x_{S_2} \cdots x_{S_k} \text{ is a basic coloring monomial}\}.$$

We call this complex *the coloring complex of G* and denote it by Δ_G.

One of the fundamental facts in the theory of face rings is that the h-vector of a d-dimensional complex Δ is given by the coefficients of the numerator when the Hilbert series of the face ring S of Δ is written as a rational function with denominator $(1-t)^{d+1}$. Thus, one customarily writes the Hilbert series $F(S,t)$ of such a ring in the form

$$F(S,t) = \frac{h(S,t)}{(1-t)^{d+1}} = \frac{h_0 + h_1 t + \cdots + h_{d+1}t^{d+1}}{(1-t)^{d+1}},$$

where $(h_0, h_1, \ldots, h_{d+1})$ is the h-vector of Δ.

The following is easy to prove.

Lemma 11 *If M is an ideal in a ring S then $F(S/M,t) = F(S,t) - F(M,t)$.* □

It is shown in Theorem 14 that the complex Δ_G, whose face ring is R/K_G, has codimension one as a subcomplex of the order complex of the Boolean algebra on d atoms, whose face ring is R. Thus, letting $S = R$ and $M = K_G$ in Lemma 11, we obtain

$$\frac{h(R/K_G,t)}{(1-t)^d} = F(R/K_G,t) = F(R,t) - F(K_G,t). \tag{1}$$

Also, the ring R is the face ring of a cone over the first barycentric subdivision of a $(d-1)$-simplex and it is well-known that its Hilbert series is given by

$$F(R,t) = \frac{1}{t} \cdot \frac{A_d(t)}{(1-t)^{d+1}},$$

where $A_d(t)$ is the d-th Eulerian polynomial, which satisfies

$$\sum_{n \geq 0} n^d t^n = \frac{A_d(t)}{(1-t)^{d+1}}.$$

Using this, together with identity (1) and Corollary 10, allows us to relate the Hilbert series of R/K_G to the *tail* of the chromatic polynomial of G, which we now define.

Definition 12 The *tail* of the chromatic polynomial χ_G of G is $T_G(n) = n^d - \chi_G(n)$.

Theorem 13 *We have*

$$\frac{1}{t} \sum_{n \geq 0} T_G(n) t^n = \frac{h(R/K_G,t)}{(1-t)^d}.$$

Thus, up to a shift by one, the W-transform of the tail T_G of χ_G equals the polynomial whose coefficients are the coordinates of the h-vector of the coloring complex of G .

Proof: We have

$$\begin{aligned} \frac{1}{t} \sum_{n \geq 0} T_G(n) t^n &= \sum_{n \geq 0} (n+1)^d t^n - \frac{1}{t} \sum_{n \geq 0} \chi_G(n) t^n = \frac{1}{t} \cdot \frac{A_d(t)}{(1-t)^{d+1}} - \frac{1}{t} \cdot \frac{W_G(t)}{(1-t)^{d+1}} \\ &= F(R,t) - F(K_G,t) = F(R/K_G,t) = \frac{h(R/K,t)}{(1-t)^d}. \end{aligned}$$ □

We shall now describe the structure of the complex Δ_G.

First, however, to facilitate the following discussion we will let the ring R be the face ring of the order complex of the *truncated Boolean algebra* $\tilde{B}_d$ on d atoms, where $\tilde{B}_d$ is B_d with $\emptyset$ and $[d]$ removed. This is a harmless modification with respect to our previous results (except in the rather trivial cases when G has fewer than three vertices) because the indeterminates $x_{[d]}$ and $x_\emptyset$ divide none of the basic coloring monomials. Thus, the Hilbert series of R/K_G is changed only in that the denominator is divided by $(1-t)^2$. The order complex of $\tilde{B}_d$ is isomorphic to the first barycentric subdivision of the boundary of a $(d-1)$-simplex, which has dimension $d-2$.

This, of course, amounts to a redefinition of the coloring complex, but even here the modification is trivial. Namely, in the complex Δ_G, $\emptyset$ and $[d]$ are both *cone points* that is, they belong to every facet of Δ_G. It is easy to show that removing a cone point (and all faces containing it) from a complex changes the h-vector only by removing its last coordinate, which is necessarily 0. Thus, the h-vector of Δ_G remains essentially the same after removing $\emptyset$ and $[d]$.

To each edge $e = ij$ of G, with $i < j$, we associate the $(d-1)!$ permutations of the letters in $S_e = [d] \setminus \{i, j\} \cup \{e\}$. We call the permutations of S_e e-permutations and we shall show that the facets of Δ_G correspond precisely to the edge permutations for G. It is important to note that, since we have removed the cone point $[d]$ from our complex, the facet corresponding to an edge-permutation does not contain the vertex $[d]$. As an example, the facet corresponding to the 25-permutation $3 - 25 - 1 - 4$ has vertices $\{3\}$, $\{2, 3, 5\}$ and $\{1, 2, 3, 5\}$.

Theorem 14 *Let G be a graph with $d \geq 3$ vertices.*

1. *To each edge of G there correspond exactly $(d-1)!$ facets of Δ_G and these are all the facets of Δ_G. The sets of such facets for two distinct edges of G are disjoint. The facets thus corresponding to an edge form a $(d-3)$-sphere, which we call an* edge-sphere *and which is isomorphic to the order complex of a truncated Boolean algebra on $(d-1)$ elements. That is, an edge-sphere is isomorphic to the first barycentric subdivision of the boundary of a $(d-3)$-simplex. In particular, Δ_G has dimension $d-3$, unless G is the graph with no edges, in which case Δ_G is the empty complex.*

2. *Any two edge-spheres intersect in a $(d-4)$-sphere which is isomorphic to the order complex of a truncated Boolean algebra on $(d-2)$ elements. Moreover, if e and $f = ij$ are two edges, then the intersection of their two spheres separates the e-sphere into two halves, where one contains all vertices of Δ_G that contain i and not j, whereas the other half contains those vertices that contain j and not i.*

Proof:

1. If G has no edges, then the permutation $\pi = d\,(d-1)\ldots 2\,1$ has no cuts except 0, since $\ell(i) = 0$ for all i. Thus, π corresponds to the empty monomial, or 1, and therefore the ideal K_G is the entire ring R, so Δ_G is the empty complex.

Suppose then that $e = ij$ is an edge in G. Let $\pi = a_1 a_2 \cdots a_d$ where, for some k, we have $a_k = i$ and $a_{k+1} = j$. Let $S_0 = \emptyset$ and let $F = \{S_1, S_2, \ldots, S_{d-2}\}$, where

$$S_m = \begin{cases} S_{m-1} \cup \{a_m\}, & \text{if } m < k, \\ S_{m-1} \cup \{i, j\}, & \text{if } m = k, \\ S_{m-1} \cup \{a_{m+1}\}, & \text{if } m > k. \end{cases}$$

Clearly, F is a facet, since it has $d-2$ vertices and thus has maximal dimension. We claim that Δ_G contains F. If Δ_G doesn't contain F then the ideal K_G must contain a monomial dividing $x_{S_1} \cdot x_{S_2} \cdot \cdots \cdot x_{S_{d-2}}$ and thus there must be a permutation with a short G-sequence of which $S_1, S_2 \setminus S_1, \ldots, S_{d-2} \setminus S_{d-3}$ is a refinement. But such a sequence must contain a set containing both i and j, which is a contradiction since ij is an edge in G (see Lemma 4). We have thus exhibited $(d-1)!$ facets F associated to the edge ij. It is easy to see that two facets thus corresponding to different edges are different. Namely, each vertex in such a facet is a set of vertices from G that either contains both or neither of a unique pair of vertices in G, and this pair of vertices constitutes the edge associated to the facet.

Conversely, we need to show that any face of Δ_G belongs to a facet associated to some edge of G as described above. Let $F = \{S_1 \subseteq S_2 \subseteq \cdots \subseteq S_k\}$ be a face of Δ_G.

We first show that some of the difference sets $D_i = S_i \setminus S_{i-1}$ must contain the vertices of an edge in G. If that is not the case then all the D_i are stable sets in G. Construct a permutation π in $\mathcal{S}_d$ by first writing all the elements of D_1 in decreasing order, then those of D_2 in order of decreasing path lengths (w.r.t. to the vertices in D_1) and in decreasing order when two vertices have the same path length associated to them (see the proof of Theorem 6). Continue this way with all the D_i's. Then $D_1, D_2, \cdots, D_k$ is a refinement of the short G-sequence of π, so the coloring ideal of G contains a monomial dividing $x_{S_1} x_{S_2} \cdots x_{S_k}$. This implies that F is a non-face of Δ_G, a contradiction, so some of the D_i's must contain an edge e of G.

This means that we can refine the chain of vertices of F down to singletons except for having one of the sets in the refinement consist of the two vertices of e. We have shown above that this refinement is a facet of Δ_G, and it is easy to see that it contains the face F.

2. If the edges $e = ij$ and $f = km$ are disjoint, then the intersection of the e-sphere and the f-sphere consists of the subcomplex of Δ_G whose vertices contain either both or neither of the vertices of the edge e and, independently, either both or neither of the vertices of the edge f. This subcomplex contains all faces of Δ_G corresponding to permutations of $[d]$ where i and j are adjacent and in increasing order and where the same is true of m and k.

   If $e = ij$ and $f = kj$ are distinct edges then the intersection of the e-sphere and the f-sphere consists of the subcomplex of Δ_G whose vertices contain either all of i, j, k or none of them. This subcomplex contains all faces of Δ_G corresponding to permutations of $[d]$ where i, j and k are three successive letters and in increasing order.

In either case it is easy to see that the subcomplex of the intersection has dimension $d-4$ and that it separates the e-sphere as claimed. Namely, there is a path along edges of Δ_G, not crossing the subcomplex, between any pair of vertices that belong to the same one of the halves described, but no such path between vertices in different halves. $\square$

Remark 15 It is possible to consider the complex obtained in the same way as the coloring complex except that we don't remove the cone point $[d]$ from Δ_G. This corresponds to associating the basic coloring monomials to the G-sequences of the permutations in $\mathcal{S}_d$ rather than their short G-sequences (and not stripping the face ring of the indeterminate $x_{[d]}$). The h-vector of this complex is the d-th Eulerian vector (coefficients of the d-th Eulerian polynomial) plus the h-vector of the coloring complex shifted one step right. Thus, knowing the h-vector of this complex is equivalent to knowing the h -vector of Δ_G. As an example, for the graphs in Figure 1 we get h-vector

$$(1, 11, 11, 1) + (0, 1, 10, 7) = (1, 12, 21, 8),$$

since the fourth Eulerian polynomial is $A_4(t) = t + 11t^2 + 11t^3 + t^4$.

Clearly, if the coloring complexes of two graphs are isomorphic, then the graphs must be *chromatically equivalent*, that is, they must have the same chromatic polynomial. However, it is possible for two non-isomorphic graphs to have isomorphic coloring complexes. Namely, there are two non-isomorphic graphs on $n \geq 4$ vertices and two edges. It follows from Theorem 14 that the coloring complexes of these graphs must be isomorphic, because each consists of two edge-spheres that intersect in a way independent of whether the edges in question are disjoint. Although these are the only examples we know of non-isomorphic graphs with isomorphic coloring complexes we do not know what the situation is in general. In Figure 1 we give the coloring complexes of the two non-isomorphic — but chromatically equivalent — graphs on three edges and four vertices. As can be seen, these complexes are not isomorphic (one has a "triangle" and the other one doesn't).

Perhaps more interesting is that the coloring complex distinguishes the two graphs given in Stanley's paper [10, Figure 1], which his chromatic symmetric function does not distinguish (see Figure 2). This can be seen as follows: A complex on a given vertex set is determined by (in fact equivalent to) its set of minimal non-faces, which in turn is equivalent to the (unique) minimal set of generators of the ideal defining its face ring. Suppose the graphs G and H in Figure 2 have isomorphic coloring complexes. Then there is a bijection ϕ between their vertex sets so that A is a minimal non-face of Δ_G if and only if $\phi(A)$ is a minimal non-face of Δ_H. Thus, the minimal sets of generators for the coloring ideals of G and H must have the

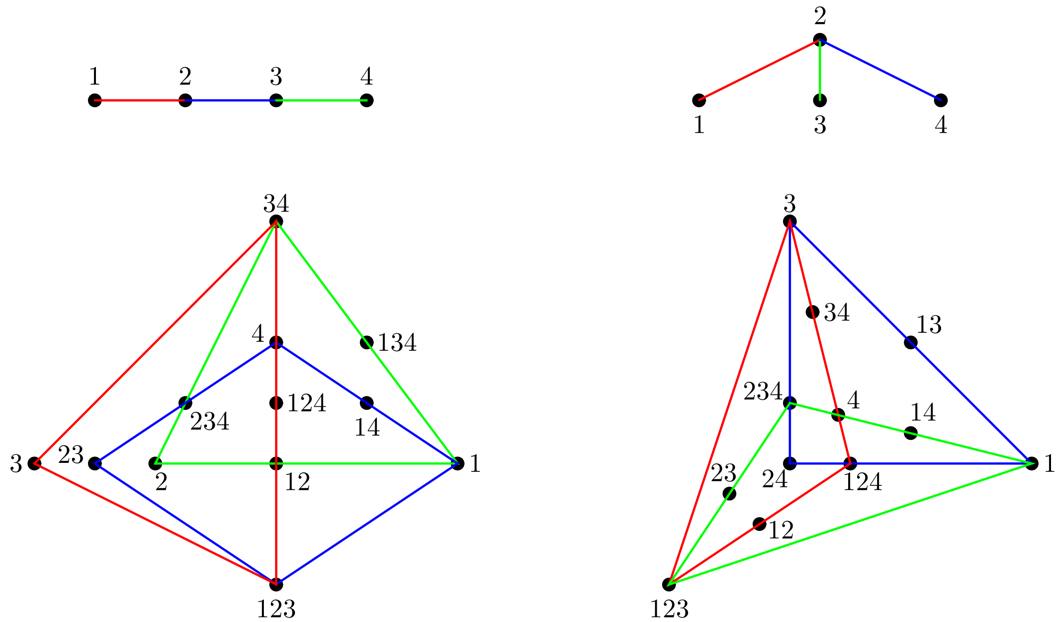

Figure 1: Two non-isomorphic graphs (trees) with the same chromatic polynomial but non-isomorphic coloring complexes.

same number of monomials of each degree. Also, the multiplicities of corresponding indeterminates in the sets of monomials constituting the respective minimal generating sets for each coloring ideal must be the same. This is not the case for the coloring ideals of G and H, which we have verified with the aid of the computer algebra program Macaulay 2 [7].

The next corollary follows from part 1 of Theorem 14. It can also be proved directly from the well-known fact that the coefficient to $-n^{d-1}$ in χ_G, which is the leading coefficient of T_G, equals the number of edges in G.

Corollary 16 *The number of facets of Δ_G, and thus the sum of the coefficients of the h-vector of Δ_G, is $E \cdot (d-1)!$, where E is the number of edges in G.* $\square$

Theorem 17 *The Euler characteristic of Δ_G equals the number of acyclic orientations of G.*

Proof: Up to a sign, the *reduced* Euler characteristic of a $(d-1)$-dimensional complex Δ is equal to the d-th coordinate $h_d(\Delta)$ of the h-vector of Δ and thus the reduced Euler characteristic of Δ_G equals the leading coefficient of the W-transform of the tail $T_G(n)$ of the chromatic polynomial. It is well known (and easy

Figure 2: Another two non-isomorphic graphs with the same chromatic polynomial and same symmetric chromatic function but non-isomorphic coloring complexes. (See [10, Figure 1].)

to prove) that the leading coefficient of the W-transform of a polynomial $P(x) = a_0 + a_1 x + \cdots + a_d x^d$ equals, up to a sign, the alternating sum of the coefficients of P. More precisely, it equals $(-1)^d P(-1)$, where d is the degree of P. Clearly, $(-1)^{d-1} T(-1) = (-1)^d \chi_G(-1) - 1$. But, by a theorem of Stanley [8, Corollary 1.3], $(-1)^d \chi_G(-1)$ equals the number of acyclic orientations of G. Since the reduced Euler characteristic is one less than the Euler characteristic, this establishes the claim. $\square$

Open problems

An obvious question is whether Δ_G is shellable. A consequence of shellability would be that Δ_G (equivalently R) is Cohen-Macaulay In that case the h-vector of the coloring complex must be an *M-vector* (see [9]), which would put certain restrictions on the values of the tail T_G and thereby on the values of the chromatic polynomial.

For $i = 1, \ldots, d$, let $\theta_i = \sum_{|S|=i} x_S$. Then it can be shown that $\theta = \theta_1, \theta_2, \ldots, \theta_d$ is a homogeneous (linear) system of parameters for R. Is R a free $k[\theta]$-module? That is equivalent to R being Cohen-Macaulay. However, a proof of this would likely be equivalent to finding a shelling of Δ_G, and shellability of Δ_G would imply that Δ_G (equivalently R) is Cohen-Macaulay.

It might be interesting to know what the minimal set of generators is for the coloring ideal of a graph G and in particular what the size of this set is. Perhaps it is more interesting to determine this for the ideal $K \cup I$, where K is the coloring ideal of G and I is the ideal used in defining the ring R, since the face ring of the coloring complex C is given by $A/(K \cup I)$.

Acknowledgements

My thanks to Eric Babson, to Ralf Fröberg, to Richard Stanley, and to an anonymous refereee for valuable comments.

References

[1] G. Almkvist: Europ. J. Combin., to appear.

[2] F. Brenti: Expansions of chromatic polynomials and log-concavity. Trans. Amer. Math. Soc. **332** (1992), no. 2, 729–756.

[3] F. Brenti: Hilbert polynomials in combinatorics, J. Algebraic Combin. **7** (1998), no. 2, 127–156;

[4] T. Y. Chow: The path-cycle symmetric function of a digraph, Adv. Math. **118** (1996), no. 1, 71–98.

[5] T. Y. Chow: Descents, quasi-symmetric functions, Robinson-Schensted for posets and the chromatic symmetric function, J. Algebraic Combin. **10** (1999), 227–240.

[6] F. R. K. Chung and R. L. Graham: On the cover polynomial of a digraph. J. Combin. Theory Ser. B **65** (1995), no. 2, 273–290.

[7] D. Grayson and M. Stillman: Macaulay2 — a software system for algebraic geometry and commutative algebra, available at `http://www.math.uiuc.edu/Macaulay2`.

[8] R. Stanley: Acyclic orientations of graphs, Discrete Math. **5** (1973), 171-178.

[9] R. Stanley: *Combinatorics and Commutative Algebra*, 2nd ed., Birkhäuser, Boston 1996.

[10] R. P. Stanley: A symmetric function generalization of the chromatic polynomial of a graph, Adv. Math. **111** (1995), no. 1, 166–194.