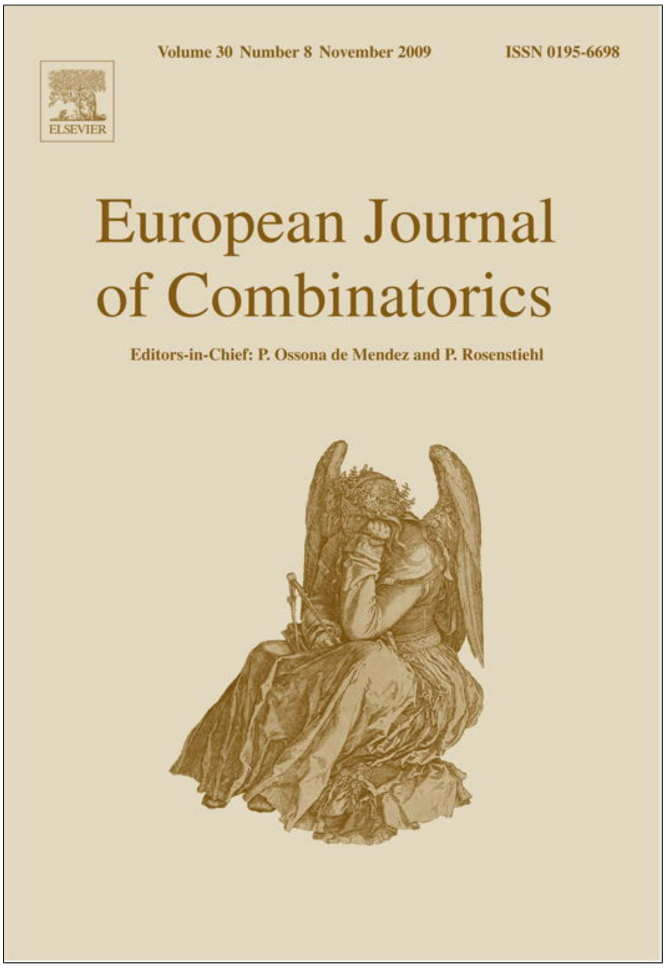
Volume 30 Number 8 November 2009
ISSN 0195-6698
ELSEVIER
European Journal of Combinatorics
Editors-in-Chief: P. Ossona de Mendez and P. Rosenstiehl

European Journal of Combinatorics 30 (2009) 1786–1800

Contents lists available at ScienceDirect

European Journal of Combinatorics

journal homepage: www.elsevier.com/locate/ejc

Monomial bases for broken circuit complexes

Jason I. Brown [a], Bruce E. Sagan [b]

[a] *Department of Mathematics and Statistics, Dalhousie University, Halifax, NS B3H 3J5, Canada*

[b] *Department of Mathematics, Michigan State University, East Lansing, MI 48824-1027, USA*

ARTICLE INFO

Article history:
Available online 21 January 2009

This paper is dedicated to Michel Las Vergnas on the occasion of his 65th birthday.

ABSTRACT

Let F be a field and let G be a finite graph with a total ordering on its edge set. Richard Stanley noted that the Stanley–Reisner ring $F(G)$ of the broken circuit complex of G is Cohen–Macaulay. Jason Brown gave an explicit description of a homogeneous system of parameters for $F(G)$ in terms of fundamental cocircuits in G. So $F(G)$ modulo this hsop is a finite dimensional vector space. We conjecture an explicit monomial basis for this vector space in terms of the circuits of G and prove that the conjecture is true for two infinite families of graphs. We also explore an application of these ideas to bounding the number of acyclic orientations of G from above.

Published by Elsevier Ltd

1. Simplicial complexes and chromatic polynomials

Let E be a finite set and let Δ be an *abstract simplicial complex on E*, i.e., a non-empty family of subsets of E such that $S \in \Delta$ and $T \subseteq S$ implies $T \in \Delta$. The elements S of Δ are called *faces*. We will assume henceforth that Δ is *pure of rank r* which means that all maximal faces S have $|S| = r$ where the absolute value sign denotes cardinality. Let $f_i = f_i(\Delta)$ be the number of $S \in \Delta$ with $|S| = i$. Then Δ has *f-vector*

$$\mathbf{f} = \mathbf{f}(\Delta) = (f_0, f_1, \dots, f_r)$$

as well as *f-polynomial*

$$f(x) = f_\Delta(x) = f_0 + f_1 x + \cdots + f_r x^r$$

where x is a variable. Henceforth we will continue the practice of appending Δ in parentheses or as a subscript when we wish to specify the complex, even if we do not do so in the corresponding definition.

E-mail addresses: brown@mscs.dal.ca (J.I. Brown), sagan@math.msu.edu (B.E. Sagan).

0195-6698/$ – see front matter. Published by Elsevier Ltd
doi:10.1016/j.ejc.2008.12.009

Another important invariant of Δ is its h-vector. Define a polynomial

$$h(x) \stackrel{\text{def}}{=} (1-x)^r f\left(\frac{x}{1-x}\right) = f_0(1-x)^r + f_1 x(1-x)^{r-1} + \cdots + f_r x^r.$$

Let h_i be the coefficient of x^i in $h(x)$ so that $h(x) = \sum_i h_i x^i$. Then the *h-vector of* Δ is

$$\mathbf{h} = (h_0, h_1, \ldots, h_r).$$

It will sometimes be convenient to extend the range of definition of the f_i and h_i by letting $f_i = h_i = 0$ if $i < 0$ or $i > r$.

Now suppose that G is a finite graph with vertices $V = V(G)$ and edges $E = E(G)$. We allow loops and multiple edges and will use the notation $p = |V|$ and $q = |E|$. We will also write $v \in G$ for $v \in V(G)$ and $e \in G$ for $e \in E(G)$ if it is clear from context whether we are talking about the vertices or edges of G. A *coloring* of G is a function $c : V \to \{1, 2, \ldots, \lambda\}$ and c is *proper* if $c(u) \neq v(v)$ for all edges $uv \in E$. Consider G's *chromatic polynomial*, $P(G) = P(G; \lambda)$, which is the number of such proper colorings. Note that if G has a loop then $P(G; \lambda) = 0$. It is well known that if G is loopless then $P(G; \lambda)$ is a monic polynomial of degree p in λ whose coefficients alternate in sign. By writing

$$P(G; \lambda) = f_0\lambda^p - f_1\lambda^{p-1} + \cdots + (-1)^p f_p \tag{1}$$

one can give the following interpretation to the coefficients f_i.

Let $\mathcal{C} = \mathcal{C}(G)$ denote the set of *cycles* of G which will also be called the set of *circuits*. Suppose G is *ordered* in that the edge set E has been given a linear ordering $e_1 < e_2 < \cdots < e_q$. Then each $C \in \mathcal{C}$ gives rise to a *broken circuit*

$$\overline{C} = C - \min C$$

where $\min C$ is the smallest edge of C in the linear ordering. The *broken complex of* G, $\Delta(G)$, is the family of all subsets of E which do not contain a broken circuit. It is easy to see that $\Delta(G)$ is a pure abstract simplicial complex. Wilf [1] was the first to consider this family of sets as a complex. In fact, $\Delta(G)$ is intimately connected with the chromatic polynomial as can be seen in the following result which dates back to Whitney [2], although he did not state it in this form.

Theorem 1.1 ([2]). *Let $P(G; \lambda)$ have coefficients f_i as defined by (1). Then*

$$f_i = f_i(\Delta(G)), \quad 0 \le i \le p. \quad \blacksquare$$

One can think of expansion (1) as being generated by a sequence of deletions and contractions expressing $P(G; \lambda)$ as a linear combination of chromatic polynomials of graphs with no edges. One could use chromatic polynomials of trees instead, or equivalently expand $P(G; \lambda)$ in terms of the basis $\{1\} \cup \{\lambda(\lambda-1)^i : i \ge 0\}$ for the ring of polynomials in λ. So define coefficients h_i by

$$P(G; \lambda) = h_0\lambda(\lambda-1)^{p-1} - h_1\lambda(\lambda-1)^{p-2} + \cdots + (-1)^p h_p. \tag{2}$$

The next result follows easily from the previous theorem and the definitions.

Corollary 1.2. *Let h_i be as in Eq. (2). Then*

$$h_i = h_i(\Delta(G)), \quad 0 \le i \le p. \quad \blacksquare$$

Our goal is to give an explicit combinatorial description of the h_i directly in terms of the broken circuits of the graph. In order to do this, we will need some machinery from the theory of Cohen–Macaulay rings.

2. Cohen–Macaulay rings and monomial ideals

Consider the polynomial ring $F[\mathbf{x}] = F[x_1, x_2, \ldots, x_q]$ where F is a field and $\mathbf{x} = \{x_1, x_2, \ldots, x_q\}$ is a set of variables. If $E = \{e_1, e_2, \ldots, e_q\}$ then any $S \subseteq E$ has corresponding monomial

$$\mathbf{x}^S = \prod_{e_i \in S} x_i.$$

Now given any simplicial complex Δ on E we form its *Stanley–Reisner ring*, $F(\Delta)$, by modding out by the non-faces of Δ, i.e.,

$$F(\Delta) = F[\mathbf{x}] \Big/ \left\langle \mathbf{x}^S : S \notin \Delta \right\rangle$$

where $\langle\cdot\rangle$ denotes the ideal generated by the polynomials in the brackets. Note that since we are generating an ideal, it suffices to consider the $\mathbf{x}^S$ where S is a minimal non-face of Δ.

If G is an ordered graph, then define

$$F(G) \stackrel{\text{def}}{=} F(\Delta(G)) = F[\mathbf{x}] \Big/ \left\langle \mathbf{x}^{\bar{C}} : C \in \mathcal{C}(G) \right\rangle$$

where we identify a (broken) circuit with its edge set. This ring has a *homogeneous system of parameters (hsop) of degree one*, i.e., a set of polynomials $\theta_1, \ldots, \theta_r \in F[\mathbf{x}]$ which are homogeneous of degree one and satisfy:

1. $\theta_1, \ldots, \theta_r$ are algebraically independent, and
2. $F(G)/\langle\theta_1, \ldots, \theta_r\rangle$ is a finite dimensional vector space over F.

Brown [3] gave an explicit construction of an hsop as follows. In order to simplify arguments, we will assume for now that $F = \mathbb{Z}_2$, the integers modulo two. In the last section, we will describe how to modify these ideas so that they will work over an arbitrary field.

First note that if G has blocks (maximal subgraphs having no cut-vertices) $G_1, G_2, \ldots, G_b$, then we have the ring isomorphism

$$F(G) \cong F(G_1) \otimes F(G_2) \otimes \cdots \otimes F(G_b) \tag{3}$$

where $\otimes$ denotes tensor product. This is because any cycle of G must be entirely contained in one of its blocks. So there is no loss of generality in assuming that G is a block and, in particular, that G is connected. Let T be a spanning tree of G. For each edge $e \in T$, let T'_e and T''_e be the components of $T - e$. So e defines a *fundamental cocircuit*

$$D_e = D_e(G) = \{uv \in E(G) : u \in T'_e,\ v \in T''_e\}$$

as well as a homogeneous degree one polynomial

$$\theta_e = \sum_{e_i \in D_e} x_i. \tag{4}$$

Since this construction will be crucial, we illustrate it with an example. Consider the graph G and its spanning tree T given in Fig. 1. For simplicity we have labeled the edges $1, 2, \ldots, 7$ rather than $e_1, e_2, \ldots, e_7$. Then we have

$$\theta_4 = x_4 + x_1 + x_2,$$
$$\theta_5 = x_5 + x_1 + x_3,$$
$$\theta_6 = x_6 + x_1 + x_2 + x_3,$$
$$\theta_7 = x_7 + x_3.$$

For any graph G, we have the following result.

Theorem 2.1 ([3]). *If G is a connected graph and T a spanning tree then the set of polynomials defined by (4) for $e \in T$ is an hsop for $\mathbb{Z}_2(G)$.* ■

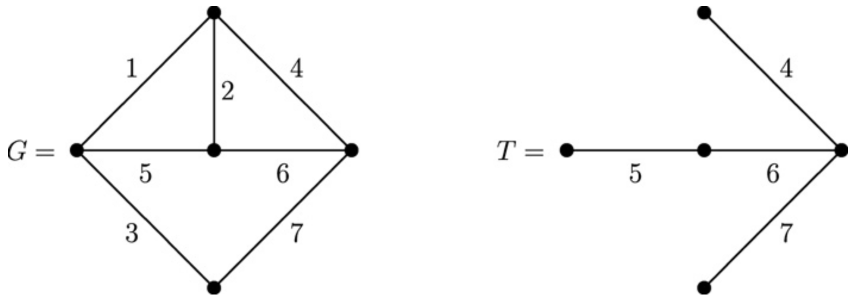

Fig. 1. A graph G and spanning tree T.

Continuing with the general development, let $\mathrm{Mon}(\mathbf{x}) = \mathrm{Mon}(q)$ denote the set of monomials in $F[\mathbf{x}] = F[x_1, x_2, \ldots, x_q]$. We will sometimes not distinguish between these monomials considered as elements of either $F[\mathbf{x}]$ or some quotient of the polynomial ring. A subset $L \subseteq \mathrm{Mon}(q)$ is a *lower order ideal* (or *down set*) if whenever $m \in L$ and $n \in \mathrm{Mon}(q)$ divides m, then $n \in L$. Similarly, $U \subseteq \mathrm{Mon}(q)$ is an *upper order ideal* (or *filter*) if whenever $m \in L$ and $n \in \mathrm{Mon}(q)$ is divisible by m, then $n \in L$. Note that U is an upper order ideal if and only if $\mathrm{Mon}(q) - U$ is a lower order ideal. If $S \subseteq \mathrm{Mon}(q)$ then the *lower and upper order ideals generated by S* are

$$L(S) = \{n \in \mathrm{Mon}(q) : n \text{ divides } m \text{ for some } m \in S\},$$
$$U(S) = \{n \in \mathrm{Mon}(q) : n \text{ is divisible by } m \text{ for some } m \in S\}.$$

Macaulay [4] showed that after modding out by an hsop, one can always find a basis of monomials which forms a lower order ideal. Also, Stanley [5] connected such a basis with the h-vector.

Theorem 2.2 ([4,5]). *Suppose that I is an ideal of $F[\mathbf{x}]$ and that $\theta_1, \ldots, \theta_r$ form an hsop for $F[\mathbf{x}]/I$. Then the ring*

$$R = \frac{F[\mathbf{x}]}{I + \langle \theta_1, \ldots, \theta_r \rangle}$$

has a basis L which is a lower order ideal of monomials.

Suppose further that $F[\mathbf{x}]/I$ is Cohen–Macaulay and $F[\mathbf{x}]/I \cong F(\Delta)$ for some simplicial complex Δ with h-vector $\mathbf{h} = (h_0, \ldots, h_r)$. Then

$$h_i = \text{number of monomials of total degree } i \text{ in } L. \quad \blacksquare$$

Now consider a graph G with a spanning tree T and define $I(G)$ to be the ideal of $F[\mathbf{x}]$ generated by the monomials $\mathbf{x}^{\overline{C}}$ for $C \in \mathcal{C}(G)$. We wish to give an explicit basis for the ring

$$R(G) = \frac{F[\mathbf{x}]}{I(G) + \langle \theta_e : e \in T \rangle}$$

which is a lower order ideal of monomials. First, however, we wish to show that we have a basis inside $\mathrm{Mon}(\mathbf{y})$ for a subset $\mathbf{y}$ of $\mathbf{x}$.

An ordering $e_1 < e_2 < \cdots < e_q$ of $E(G)$ will be called *standard* if the last $p - 1$ edges in the order form a tree. From now on we will assume that all our orderings are standard and take our spanning tree $T = T(G)$ to be the one determined by the last edges in the order. It will also be convenient to denote the number of edges not in T by $k = q - p + 1$. We will show that our basis can be taken in $\mathrm{Mon}(\mathbf{y})$ where $\mathbf{y} = \{x_1, x_2, \ldots, x_k\}$.

We now return to working over $\mathbb{Z}_2$. Suppose $k < j \leq q$ and write D_j for D_{e_j} and θ_j for θ_{e_j}. Then since $\theta_j = 0$ in $R(G)$ we have

$$x_j = \sum_{e_i \in (D_j - e_j)} x_i \tag{5}$$

where $x_i \in \mathbf{y}$ for all x_i appearing in the sum. For each $C \in \mathcal{C}$ let $p_{\overline{C}} = p_{\overline{C}}(\mathbf{y})$ be the polynomial obtained from $x^{\overline{C}}$ by substituting in the sum in Eq. (5) for x_j for each $j > k$. Consider the ideal

$$J = J(G) = \langle p_{\overline{C}} : C \in \mathcal{C} \rangle.$$

We immediately have the following result.

Proposition 2.3. *If G is a connected graph and $F = \mathbb{Z}_2$ then*

$$R(G) \cong \frac{\mathbb{Z}_2[\mathbf{y}]}{J(G)}. \quad \blacksquare$$

Returning to our running example, we convert the list of circuits in G into polynomials using the equations for $\theta_4, \ldots, \theta_7$.

$$\begin{array}{lll}
C_1 = \{1,4,5,6\}, & \mathbf{x}^{\overline{C_1}} = x_4x_5x_6, & p_{\overline{C_1}} = (x_1+x_2)(x_1+x_3)(x_1+x_2+x_3),\\
C_2 = \{2,4,6\} & \mathbf{x}^{\overline{C_2}} = x_4x_6, & p_{\overline{C_2}} = (x_1+x_2)(x_1+x_2+x_3),\\
C_3 = \{3,5,6,7\} & \mathbf{x}^{\overline{C_3}} = x_5x_6x_7, & p_{\overline{C_3}} = x_3(x_1+x_3)(x_1+x_2+x_3),\\
C_4 = \{1,2,5\} & \mathbf{x}^{\overline{C_4}} = x_2x_5, & p_{\overline{C_4}} = x_2(x_1+x_3),\\
C_5 = \{1,3,4,7\} & \mathbf{x}^{\overline{C_5}} = x_3x_4x_7, & p_{\overline{C_5}} = x_3^2(x_1+x_2),\\
C_6 = \{2,3,4,5,7\} & \mathbf{x}^{\overline{C_6}} = x_3x_4x_5x_7, & p_{\overline{C_6}} = x_3^2(x_1+x_2)(x_1+x_3),\\
C_7 = \{1,2,3,6,7\} & \mathbf{x}^{\overline{C_7}} = x_2x_3x_6x_7, & p_{\overline{C_7}} = x_2x_3^2(x_1+x_2+x_3).
\end{array}$$

We will now pick a specific monomial $m_{\overline{C}}$ from each $p_{\overline{C}}$ and these will be used to define the lower order ideal of monomials being sought. For $1 \le i \le k$, the graph $T + e_i$ has a unique circuit C_i and these circuits will be called *fundamental*. We label the non-fundamental circuits in some order as C_i for $i > k$. Also define

$$d_i = \begin{cases} i & \text{if } i \le k,\\ \min\{j : e_j \in D_i\} & \text{if } i > k.\end{cases}$$

Now let

$$m_{\overline{C}_i} = \begin{cases} x_i^{|\overline{C}_i|} & \text{if } i \le k,\\ \prod_{e_j \in \overline{C}_i} x_{d_j} & \text{if } i > k.\end{cases}$$

It is easy to see from the definitions that $m_{\overline{C}}$ is indeed a term in the polynomial $p_{\overline{C}}$. Finally, define upper and lower order ideals

$$U(G) = U(m_{\overline{C}} : C \in \mathcal{C}(G)) \quad \text{and} \quad L(G) = \mathrm{Mon}(k) - U(G).$$

Note that all these quantities depend on the ordering imposed on the edges and not just on the graph itself, even though our notation does not reflect that. It is $L(G)$ which will be our candidate as a monomial basis for $R(G)$.

Continuing with our example, C_1, C_2, and C_3 are fundamental with 4, 3, and 4 edges (respectively) and so

$$m_{\overline{C}_1} = x_1^3, \qquad m_{\overline{C}_2} = x_2^2, \qquad m_{\overline{C}_3} = x_3^3.$$

The monomials $m_{\overline{C}}$ for the other four circuits are obtained by taking the variable of smallest subscript in each factor of the corresponding $p_{\overline{C}}$, so

$$m_{\overline{C}_4} = x_1x_2, \qquad m_{\overline{C}_5} = x_1x_3^2, \qquad m_{\overline{C}_6} = x_1^2x_3^2, \qquad m_{\overline{C}_6} = x_1x_2x_3^2.$$

Thus $R(G)$ should have as basis

$$\begin{aligned} L(G) &= \mathrm{Mon}(3) - U(x_1^3, x_2^2, x_3^3, x_1x_2, x_1x_3^2, x_1^2x_3^2, x_1x_2x_3^2)\\ &= \{1, x_1, x_2, x_3, x_1^2, x_3^2, x_1x_3, x_2x_3, x_1^2x_3, x_2x_3^2\}\end{aligned}$$

and this can be verified directly.

A graph for which there is an ordering of E such that $L(G)$ is a basis for $R(G)$ will be said to have a *no broken circuit basis* or *NBC basis*. The rest of this paper is organized as follows. In the next section we will prove a general theorem about when a block has an NBC basis. In Section 4 we will apply these ideas to show that two infinite families of graphs do indeed have NBC bases. Section 5 will be devoted to giving an upper bound for the number of acyclic orientations for a graph with an NBC basis. We also compare this bound to others in the literature. We end with some comments and open problems, including a possible approach to proving our conjecture that every graph G has an ordering which produces an NBC basis for $R(G)$.

3. Graphs with NBC bases

In order to show that a graph G has an NBC basis, we could try induction on the number of edges. Since the chromatic polynomial is involved, this would entail deletion and contraction. If $e \in E(G)$ then let $G\backslash e$ and G/e denote G with e deleted and with e contracted, respectively. Since loops and multiedges are allowed, both $G\backslash e$ and G/e will have exactly one less edge than G. An elementary fact about the chromatic polynomial is that

$$P(G; \lambda) = P(G\backslash e; \lambda) - P(G/e; \lambda).$$

Using this equation and (2) we easily obtain the following proposition.

Proposition 3.1. *Let G be a graph and $e \in E(G)$. Then for all $i \geq 0$ we have*

$$h_i(G) = h_i(G\backslash e) + h_{i-1}(G/e). \quad \blacksquare$$

If we choose $e \in T$ then T/e is a spanning tree of G/e but $T\backslash e$ is no longer a tree. On the other hand, if we choose $e \notin T$ then T is still a spanning tree of $G\backslash e$ but T is no longer a tree in G/e. However, we can avoid these difficulties if G has a vertex w with $\deg w = 2$ where $\deg w$, the *degree of* w, is the number of edges containing w.

As noted before, we can restrict our attention to graphs G which are blocks so that $G\backslash e$ and G/e are connected for all $e \in E$. We will say that a standard ordering $e_1 < e_2 < \cdots < e_q$ on G imposes the *induced ordering* $e_1 < e_2 < e_3 < \cdots < e_{q-1}$ on $G\backslash e_q$ and on G/e_q. Now suppose that G has a vertex w with $\deg w = 2$ and that e_k, e_q are the two edges containing w. If the ordering on G is standard, then so are the induced orderings on $G\backslash e_q$ and G/e_q. Our primary tool for showing that certain graphs have NBC bases will be the following theorem. Note that an example which illustrates the proof of this result follows the demonstration, so the reader may wish to read both in parallel.

Theorem 3.2. *Let G be a block with a standard ordering $e_1 < e_2 < \cdots < e_q$. Suppose G has a vertex w of degree two such that the edges containing w are e_k and e_q. If $R(G\backslash e_q)$ and $R(G/e_q)$ have NBC bases in their induced standard orderings, then so does $R(G)$.*

Proof. Let $\uplus$ denote disjoint union, and if $S \subseteq \mathrm{Mon}(k)$ and $m \in \mathrm{Mon}(k)$ then let $mS = \{mn : n \in S\}$. We will first show that

$$L(G) = L(G\backslash e_q) \uplus x_k L(G/e_q). \tag{6}$$

So by our assumptions about $R(G\backslash e_q)$ and $R(G/e_q)$ and the previous proposition (summed over all i), we have

$$|L(G)| = |L(G\backslash e_q)| + |L(G/e_q)| = \dim R(G\backslash e_q) + \dim R(G/e_q) = \dim R(G)$$

where dimension is being taken over the field $\mathbb{Z}_2$.

Consider $G\backslash e_q$. Note that e_k is in the tree for $G\backslash e_q$ and so the basis for $R(G\backslash e_q)$ will be in $\mathrm{Mon}(k-1)$. Also, from our assumptions on w, e_k is the only edge of $G\backslash e_q$ containing w. So C is a circuit of $G\backslash e_q$ if and only if C is a circuit of G not containing e_k. It follows that x_k is never a factor of $\mathbf{x}^{\overline{C}}$ for such C. It also follows that for $e_j \in T(G\backslash e_q)$, $e_j \neq e_k$, we have $D_j(G\backslash e_q) = D_j(G) - e_k$. And both of these sets have the same minimum since e_k is the edge of largest index outside the tree for G. Thus the generators for

$J(G\backslash e_q)$ are obtained from those for $J(G)$ by setting $x_k = 0$ wherever it appears. So the monomials in $U(G\backslash e_q)$ are precisely those in $U(G)$ which do not have x_k as a factor. Hence $L(G\backslash e_q)$ consists of the monomials in $L(G)$ which do not have x_k as a factor.

Now consider G/e_q. The circuits of G are in bijection with the circuits of G/e_q: If $C \in \mathcal{C}(G)$ contains e_q then it corresponds to the circuit C/e_q of G/e_q, while if C does not contain e_q then it is also a circuit of G/e_q itself. We will call the former circuits (in both G and G/e_q) *type I*, and the latter *type II*. Note that because of the assumptions on w, the type I and type II circuits can also be characterized as those which do and do not contain e_k, respectively. Since e_q is the only edge of $T(G)$ containing w, we have $D_j(G/e_q) = D_j(G)$ for each $e_j \in T(G/e_q)$. Thus, using $\tilde{p}_{\overline{C}}$ to denote the generators of $J(G/e_q)$,

$$p_{\overline{C}} = \begin{cases} x_k \tilde{p}_{\overline{C/e_q}} & \text{if } C \text{ is of type I,} \\ \tilde{p}_{\overline{C}} & \text{if } C \text{ is of type II,} \end{cases}$$

where the polynomials for the type II circuits have no factor of x_k. Since e_k has the largest index outside $T(G)$, the same relation holds between the corresponding generators of $U(G)$ and of $U(G/e_q)$, i.e., $m_{\overline{C}} = x_k \tilde{m}_{\overline{C/e_q}}$ or $\tilde{m}_{\overline{C}}$ depending on whether $\overline{C}$ is type I or type II (respectively), where the tilde indicates the quantity is being calculated in C/e_q.

We claim that $x_k L(G/e_q)$ consists precisely of the monomials in $L(G)$ which have a factor of x_k: Suppose that we have a monomial of $L(G)$ divisible by x_k. Then it can be written as $x_k m$ for some $m \in \text{Mon}(k)$. Since $x_k m$ is not divisible by any type I generator of $U(G)$, and all such generators have the form $x\tilde{m}_1$ for some type I generator $\tilde{m}_1$ of $U(G/e_q)$, we see that m is not divisible by $\tilde{m}_1$ for all type I generators of $U(G/e_q)$. Also, $x_k m$ is not divisible by any type II generator $\tilde{m}_2$ of $U(G)$, and all such generators do not have x_k as a factor, so m is not divisible by any type II generator $\tilde{m}_2$ of $U(G/e_q)$. So $m \in L(G/e_q)$ and $x_k m \in x_k L(G/e_q)$. The proof of the reverse inclusion is similar.

Since $L(G)$ is clearly the disjoint union of its monomials with a factor of x_k and its monomials without a factor of x_k, we are done with the demonstration of (6). So we have proved that $L(G)$ contains $\dim R(G)$ monomials, and thus it will suffice to show that these monomials span $R(G)$. For that, it suffices to show that $L(G)$ spans $U(G)$. So take $m \in U(G)$. Suppose first that x_k is a factor of m, that is, $m = x_k n$ for some monomial n. Then from the previous paragraph we have $n \in U(G/e_q)$. So by our assumption about $R(G/e_q)$, we can write

$$n = \sum_{l \in L(G/e_q)} a_l l + p \tag{7}$$

where the a_l are constants and $p \in J(G/e_q)$. But $x_k l \in L(G)$ for $l \in L(G/e_q)$, and $x_k p \in J(G)$ for $p \in J(G/e_q)$ since this is true for each of the generators of $J(G/e_q)$. So multiplying (7) by x_k expresses $m = x_k n$ as a linear combination of elements of $L(G)$ modulo $J(G)$ as desired.

Now suppose that x_k is not a factor of m. Then by our previous results concerning $G\backslash e_q$ we have $m \in U(G\backslash e_q)$. Thus, by our assumption about $R(G\backslash e_q)$, we can write

$$m = \sum_{l \in L(G\backslash e_q)} a_l l + p \tag{8}$$

where the a_l are constants and $p \in J(G\backslash e_q)$. Now, as shown above, $l \in L(G\backslash e_q)$ implies $l \in L(G)$. Furthermore, there must be a $p' \in J(G)$ such that $p'(x_1, \ldots, x_{k-1}, 0) = p$. It follows that $p' = p + x_k p''$ for some $p'' \in F[\mathbf{y}]$. But, from the previous paragraph, we have that $x_k p''$ is spanned by $L(G)$ modulo $J(G)$. So substituting $p = p' - x_k p''$ into (8) expresses m as a linear combination of elements of $L(G)$ modulo $J(G)$. Hence every monomial is in the span of $L(G)$ as required. ■

Returning to our example graph (which satisfies the conditions of the previous theorem), $G\backslash e_7$ and the tree for the induced order are shown in Fig. 2. The relevant sets are

$$\{\theta_e\} = \{x_3,\ x_4 + x_1 + x_2,\ x_5 + x_1,\ x_6 + x_1 + x_2\},$$
$$\{C\} = \{\{1, 4, 5, 6\},\ \{2, 4, 6\},\ \{1, 2, 5\}\},$$
$$\{\mathbf{x}^{\overline{C}}\} = \{x_4 x_5 x_6,\ x_4 x_6,\ x_2 x_5\},$$

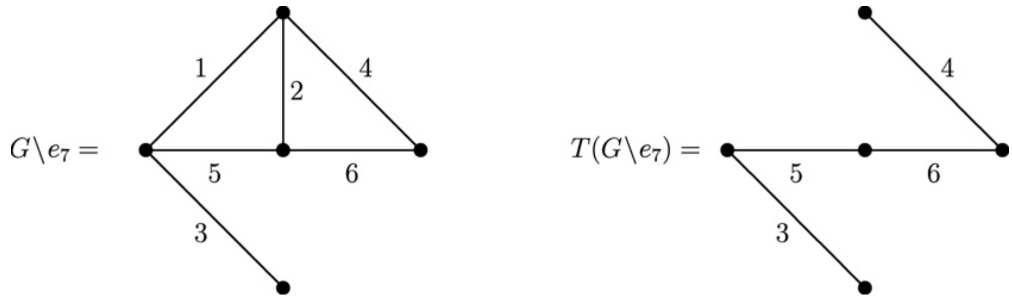

Fig. 2. The graph $G\backslash e_7$ and spanning tree $T(G\backslash e_7)$.

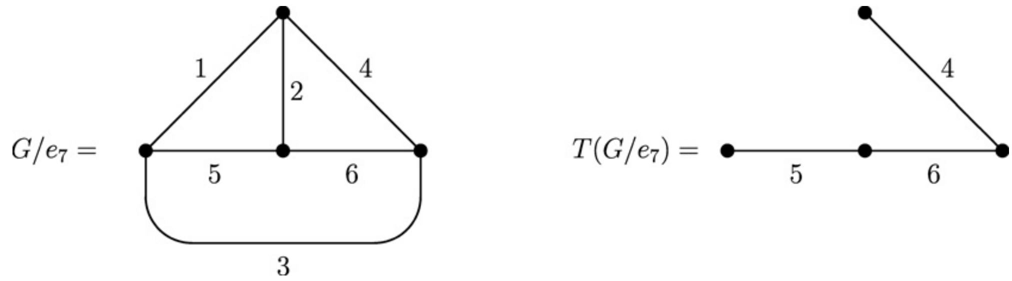

Fig. 3. The graph G/e_7 and spanning tree $T(G/e_7)$.

$$\{p_{\overline{C}}\} = \{x_1(x_1+x_2)^2,\ (x_1+x_2)^2,\ x_1x_2\},$$
$$U(G\backslash e_7) = U(x_1^3,\ x_2^2,\ x_1x_2),$$
$$L(G\backslash e_7) = Mon(2) - U(G\backslash e_7) = \{1,\ x_1,\ x_2,\ x_1^2\}.$$

Making the same computations in G/e_7 (Fig. 3) yields

$$\{\theta_e\} = \{x_4+x_1+x_2,\ x_5+x_1+x_3,\ x_6+x_1+x_2+x_3\},$$
$$\{C\} = \{\{1,4,5,6\},\ \{2,4,6\},\ \{3,5,6\},\ \{1,2,5\},\ \{1,3,4\},\ \{2,3,4,5\},\ \{1,2,3,6\}\},$$
$$\{\mathbf{x}^{\overline{C}}\} = \{x_4x_5x_6,\ x_4x_6,\ x_5x_6,\ x_2x_5,\ x_3x_4,\ x_3x_4x_5,\ x_2x_3x_6\},$$
$$\{p_{\overline{C}}\} = \{(x_1+x_2)(x_1+x_3)(x_1+x_2+x_3),\ (x_1+x_2)(x_1+x_2+x_3), (x_1+x_3)(x_1+x_2+x_3),\ x_2(x_1+x_3),\ x_3(x_1+x_2),\ x_3(x_1+x_2)(x_1+x_3),\ x_2x_3(x_1+x_2+x_3)\},$$
$$U(G/e_7) = U(x_1^3,\ x_2^2,\ x_3^2,\ x_1x_2,\ x_1x_3,\ x_1^2x_3,\ x_1x_2x_3),$$
$$L(G/e_7) = \mathrm{Mon}(3) - U(G/e_7) = \{1,\ x_1,\ x_2,\ x_3,\ x_1^2,\ x_2x_3\}.$$

Note that we have $L(G) = L(G\backslash e_7) \uplus x_3L(G/e_7)$.

4. Two families

We will now consider two families of graphs and prove that they have NBC bases. They are called (generalized) theta and phi graphs.

A *(generalized) theta graph* consists of two vertices u, v together with t internally-disjoint u–v paths $P^{(1)}, P^{(2)}, \ldots, P^{(t)}$. Note that we are not insisting that $t = 3$ as is usually done for theta graphs. By convention, we assume that the paths are listed in weakly increasing order of length. To show that such a graph G has an NBC basis, we need to label its edges so that $e_1 < e_2 < \cdots < e_q$ is a standard order. We can obtain a spanning tree of G by removing an edge from all but one of the paths, so $t = k + 1$. Now for $1 \le i \le k + 1$, let e_i label the edge of $P^{(i)}$ containing u. Label the remaining edges $e_q, e_{q-1}, \ldots, e_{k+2}$ where the edges of $P^{(1)}\backslash e_1$ come first, those of $P^{(2)}\backslash e_2$ come second, and so forth, with the edges on each $P^{(i)}$ listed in order starting from the one adjacent to e_i. An example of this labeling is given in Fig. 4.

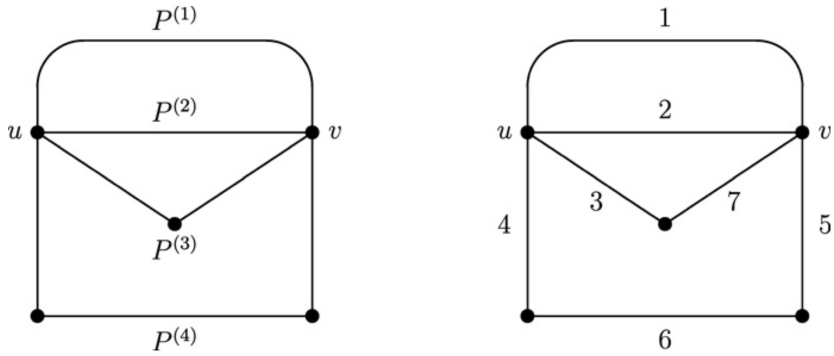

Fig. 4. A theta graph and its labeling.

Before proving that a theta graph has an NBC basis, we will need a lemma to take care of the special case when there is more than one u–v path of length 1, that is, there is a multiedge uv of multiplicity at least 2. Let G be a connected graph with standard ordering $e_1 < e_2 < e_3 < \cdots < e_q$ where, for some $l \le k$, e_l and e_{l-1} have the same endpoints. Let $G \backslash e_l$ have the induced ordering $e_1 < \cdots < \hat{e}_l < \cdots < e_q$ where the hat indicates that e_l has been removed. Note that the induced ordering is standard. Then the corresponding rings are related in the manner in which one would expect given that the chromatic polynomials do not change.

Lemma 4.1. *Suppose that G has a standard ordering such that e_l and e_{l-1} have the same endpoints for some $l \le k$. If $G \backslash e_l$ is given the induced ordering above then $R(G) \cong R(G \backslash e_l)$.*

Proof. The generators for $J(G)$ can be obtained from those for $J(G \backslash e_l)$ by replacing x_{l-1} with $x_{l-1} + x_l$ everywhere. The additional cycle made by $\{e_{l-1}, e_l\}$ gives $x_l = 0$ in the quotient $R(G)$, and the isomorphism follows. ∎

Theorem 4.2. *If G is a (generalized) theta graph with a theta labeling then G has an NBC basis.*

Proof. We will induct on the number of edges of G. From the previous lemma and our convention for listing the $P^{(i)}$, we can reduce to the case where there is at most one u–v path of length 1. Furthermore, when there are only one or two u–v paths the result is easy to verify. (In the latter case, G is just a single cycle and the details appear in [3].) So we can assume that G is a block and has at least two u–v paths of length at least 2.

Let $P^{(i)}$ be the first path on the path list for G which has length at least 2. Let w be the vertex on $P^{(i)}$ adjacent to u. Then because of our labeling method and the fact that $P^{(i+1)}$ has length at least 2, w satisfies the hypotheses of Theorem 3.2 except that w is adjacent to e_q and e_l for some $l \le k$, not necessarily e_k itself. Recall that in the proof of that theorem, the edge e_k was chosen because k was the largest index outside $T(G)$. This guaranteed that for each circuit C, the monomials $m_{\overline{C}}$ would be the same in G, $G \backslash e_q$, and G/e_q if C was type II, but would differ by a factor of x_k between G and G/e_q if C was type I. But the same statements hold in a theta graph, with x_l in place of x_k, because of its geometry and the way we have chosen the labeling. Thus we can use induction in the same manner as done in Theorem 3.2.

Now consider $G \backslash e_q$. This is not a theta graph in general. But the induced labeling on $G \backslash e_q$ is a theta labeling if we ignore the other edges on $P^{(i)}$. This does not cause any problems since each of these edges is now a block and so does not contribute anything to $F(G)$ by (3) and the fact that $R(e) \cong F$ for any edge e. Hence, by induction, $R(G \backslash e_q)$ has an NBC basis.

Now look at G/e_q. This is still a theta graph and its induced labeling is a theta labeling. So, by induction, $R(G/e_q)$ has an NBC basis. Hence, by the same reasoning as in Theorem 3.2, G has an NBC basis. ∎

As a special case of the previous result, we obtain the following.

Corollary 4.3. *The complete bipartite graph $K_{2,t}$ with a theta labeling has an NBC basis.* ∎

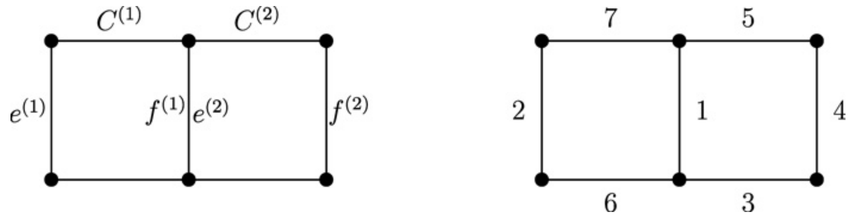

Fig. 5. The phi graph $P_2 \times P_3$ and one of its phi labelings.

Rather than thinking of theta graphs as unions of paths, one could consider them as a set of cycles joined in parallel. We will now define a family of graphs which can be thought of as joining cycles in series. Suppose we are given t cycles $C^{(1)}, C^{(2)}, \ldots, C^{(t)}$ all of length at least two, and in each $C^{(i)}$ we are given a pair of distinguished edges $e^{(i)}, f^{(i)}$. Then the associated *phi graph* is obtained by identifying $f^{(i)}$ with $e^{(i+1)}$ for $1 \leq i < t$. For example, if we let P_p denote the path on p vertices then the cross product $P_2 \times P_t$ is a phi graph where all the $C^{(i)}$ have length four. (It is because of the shape of $P_2 \times P_3$ that we call these phi graphs.)

Again, we will need a specific labeling for a phi graph G. Note that we can obtain a spanning tree for G by removing one edge from each $C^{(i)}$, so $t = k$. Label edge $e^{(i)}$ with e_{k-i+1}, $1 \leq i \leq k$. Now label the remaining edges of $C^{(1)}$ as follows. We have $C^{(1)} - e^{(1)} - f^{(1)} = P \uplus Q$ where P, Q are paths. Label the edges along P (if any) starting with the one adjacent to $e^{(1)}$ with $e_q, e_{q-1}, \ldots, e_{r+1}$. Now do the same along Q using the labels $e_r, e_{r-1}, \ldots, e_s$. Continue in a similar manner to label the rest of the cycles. (When one gets to the last one, there will be only one path to label.) Call this a *phi labeling* of the graph. The graph $P_2 \times P_3$ and one of its phi labelings are given in Fig. 5.

Theorem 4.4. *If G is a phi graph with a phi labeling then G has an NBC basis.*

Proof. Again, we induct on the number of edges in G. As mentioned previously, the case of a single cycle has already been done. So suppose we have at least two $C^{(i)}$. If $C^{(1)}$ has length 2, then its phi labeling is exactly of the type considered in Lemma 4.1. Thus $R(G) \cong R(G \backslash e_k)$ where the latter graph has a phi labeling and fewer edges, and the result follows by induction. If $C^{(1)}$ has length at least 3, then a deletion–contraction argument similar to the one used for theta graphs will provide a proof (the details are left to the reader). ∎

Corollary 4.5. *The graph $P_2 \times P_t$ with a phi labeling has an NBC basis.* ∎

5. Upper bounds

If graph $G = (V, E)$ has an NBC basis, then we can use this fact to give a simple upper bound on its h-vector. (Lower bounds for h-vectors of various types of complexes have been given by Swartz [6].) This, in turn, bounds the values of the chromatic polynomial $P(G; \lambda)$ at negative integers since then all terms in the expansion (2) have the same sign. In particular, this gives an upper bound on $\alpha(G)$, the number of acyclic orientations of G, because of a famous theorem of Stanley [7] which states that

$$\alpha(G) = (-1)^p P(G; -1)$$

where, as usual, $p = |V|$. To see why one could only expect to bound these quantities, rather than obtaining their exact values, we need to say a few words about the theory of #P problems which was introduced by Valiant [8,9].

If A and B are two problems then we say that A is *polynomially reducible* to B if it is possible, given a subroutine to solve B, to solve A in polynomial time, where we count calls to the subroutine for B as a single step. The *class #P* consists of those enumeration problems where the structures being counted can be recognized in polynomial time. In other words, there is an algorithm which is polynomial in the size of the input problem that can verify whether a given structure should be included in the count. So the class #P is for enumeration problems as the class NP is for decision problems. An enumeration problem is *#P-complete* if any problem in #P is polynomially reducible to it. So the #P-complete problems are the hardest in #P.

Linial [10] first showed that computing $\alpha(G)$ is #P-complete. Jaeger, Vertigan, and Welsh [11] derived more general results about computing the Tutte polynomial of a matroid which imply that computing $P(G;\lambda)$ is #P-complete for all but nine special values of λ.

The case $\lambda = -1$ has attracted special interest because $\log\alpha(G)$ is a lower bound on the computational complexity of certain decision and sorting problems, see for example the paper of Goddard, Kenyon, King, and Schulman [12]. Obviously the number of acyclic orientations of G is bounded above by the total number of orientations, giving

$$\alpha(G) \le 2^q$$

where $q = |E|$. Fredman (whose work is reported in a paper of Graham, Yao, and Yao [13, Section 7]), and independently Manber and Tompa [14] gave the first non-trivial upper bound for $\alpha(G)$ as

$$\alpha(G) \le \prod_{v\in V}(\deg v + 1),$$

where, as usual, $\deg v$ is the degree of vertex v. This bound was improved by Kahale and Schulman [15] as follows.

Given a graph G, consider its *cone*, G^*, obtained by adding a new vertex adjacent to every vertex of G. Then Kahale and Schulman show that $\alpha(G)$ is at most the number of spanning trees of G^*. Using the Matrix-Tree Theorem, this bound can be expressed as a determinant. Since the determinant itself could be costly to compute, they give an upper bound for its value.

Theorem 5.1 (*[15]*). *We have the upper bound*

$$\alpha(G) \le \prod_{v\in V}(\deg v + 1)\prod_{uw\in E}\exp\frac{-1}{2(\deg u + 1)(\deg w + 1)} \stackrel{\text{def}}{=} \beta(G). \quad \blacksquare \tag{9}$$

Now suppose that G has an NBC basis $\mathrm{Mon}(k) - U(m_{\overline{C}} : C \in \mathcal{C}(G))$. If we remove the upper order ideal generated by just the fundamental circuits, then we will get a spanning set for the quotient which can be used to bound the h-vector from above. Furthermore, each of these monomials has the simple form

$$m_{\overline{C}_i} = x_i^{|C_i|-1}.$$

So by Theorem 2.2 and Eq. (2), we have proved the following result, where we use $L_d(S)$ to denote the set of monomials in the lower order ideal $L(S)$ which have total degree d.

Theorem 5.2. *If G has an NBC basis with fundamental circuits $C_1, \ldots, C_k$ then, for $d \ge 0$,*

$$h_d(G) \le \left|L_d\left(x_1^{|C_1|-2}\cdots x_k^{|C_k|-2}\right)\right| \stackrel{\text{def}}{=} l_d(G). \tag{10}$$

Furthermore

$$\alpha(G) \le \sum_{d=0}^{p-1} l_d(G)2^{p-d-1} \stackrel{\text{def}}{=} \gamma(G). \quad \blacksquare \tag{11}$$

It is an easy exercise to show that

$$l_d(G) \le |L_d(\mathrm{Mon}(k))| = \binom{d+k-1}{k-1}. \tag{12}$$

One can calculate the exact values of the $l_d(G)$ using the Principle of Inclusion–Exclusion (see Stanley's text [16, Chapter 2]).

We will now compare the bounds $\beta(G)$ and $\gamma(G)$ for certain theta and phi graphs. When possible, we will compare the γ bound with the actual number of acyclic orientations. Of course, from a practical viewpoint, it is unnecessary to use a bound when the exact value is known. But this will show how close γ is to the real value.

We keep the conventions of the previous section. Define $\Theta_{n,t}$ to be the theta graph consisting of t paths of length n with their endpoints identified to form the special vertices u and v. There is an interesting change in the behaviour of the γ bound depending on whether n is fixed and t varies, or vice versa.

Theorem 5.3. *As $n \to \infty$ we have*

$$\gamma(\Theta_{n,3}) \sim \alpha(\Theta_{n,3}).$$

As $t \to \infty$ we have

$$\beta(\Theta_{2,t}) = o(\gamma(\Theta_{2,t})).$$

Proof. First consider $\Theta_{n,3}$ where $p = 3n - 1$ and $q = 3n$. Since this graph only has 3 circuits, it is easy to use Inclusion–Exclusion to calculate $\alpha(G)$, from which one sees that the count is asymptotic to the first term

$$\alpha(G) \sim 2^q = 2^{3n}.$$

To compute γ, first note that from (10) and (12) we have

$$h_d(\Theta_{n,3}) \le l_d(x_1^{2n-2}\, x_2^{2n-2}) \le d + 1.$$

Plugging this bound into (11) gives

$$\gamma(\Theta_{n,3}) \le \sum_{d\ge 0}(d+1)2^{3n-2-d} = \frac{2^{3n-2}}{(1-1/2)^2} = 2^{3n}.$$

So we must also have $\gamma(\Theta_{n,3}) \sim 2^{3n}$ since γ is an upper bound.

For $\Theta_{2,t}$ note that k, the number of edges not in a spanning tree, satisfies $k = t - 1$. We also have $p = t + 2$ and $q = 2t$. Using (10), we get

$$l(d,t) \stackrel{\text{def}}{=} l_d(\Theta_{2,t}) = \left| L_d\left(x_1^2 x_2^2 \cdots x_{t-1}^2\right)\right|$$

which is the coefficient of y^d in the expansion of the generating function $(1 + y + y^2)^{t-1}$. From this, it follows that the $l(d,t)$ satisfy the recursion

$$l(d,t+1) = l(d,t) + l(d-1,t) + l(d-2,t). \tag{13}$$

Let $\gamma_t = \gamma(\Theta_{2,t})$. So multiplying (13) by 2^{t+2-d} and summing over $0 \le d \le t+2$, we can use (11) to get the following equation, with the three expressions in brackets coming from the three terms of the recursion (respectively):

$$\gamma_{t+1} = [2\gamma_t + l(t+2,t)] + [\gamma_t] + \left[\frac{1}{2}\gamma_t - \frac{1}{2}l(t+1,t)\right] > \frac{7}{2}\gamma_t - \frac{1}{4}\gamma_t = \frac{13}{4}\gamma_t \tag{14}$$

where the inequality follows by noting $4l(t-1,t)$ is a summand in γ_t and that, as provable from the generating function, the sequence $(l(d,t))_{0\le d\le 2t-2}$ is symmetric and unimodal with maximum at $l(t-1,t)$.

Finally, combining the estimates in (9) and (14), we see that for any $0 < \epsilon < 1/4$,

$$\beta(\Theta_{2,t}) = (t+1)^2 3^t \exp\frac{-2t}{6t+6} = o((13/4-\epsilon)^t) = o(\gamma(\Theta_{2,t}))$$

as desired. ■

Now for $n \ge 4$, let $\Phi_{n,t}$ be a phi graph derived by pasting together t cycles of length n in such a way that each cycle only intersects the cycle just preceding and the cycle just following it (if any). Note that $\Phi_{n,t}$ is actually a graph family since one can get a number of graphs with these specifications by pasting along different edges. But they all have a uniform description of their NBC bases and degree sequences, so the bounds under consideration will apply to any graph of the family.

Theorem 5.4. *As $n \to \infty$ we have*

$$\gamma(\Phi_{n,2}) \sim \alpha(\Phi_{n,2}).$$

As $t \to \infty$ we have

$$\gamma(\Phi_{4,t}) = o(\beta(\Phi_{4,t})).$$

Proof. The proof for $\Phi_{n,2}$ is completely analogous to the proof given for $\Theta_{n,3}$, so we leave it to the reader.

By considering $P_2 \times P_{t+1}$ or any other member of $\Phi_{4,t}$, we see that $p = 2t + 2$, $q = 3t + 1$, and $k = t$. Using bound (12) and the Binomial Theorem in (11) yields

$$\gamma(\Phi_{4,t}) \le \sum_{d=0}^{2t+1} \binom{d+t-1}{t-1} 2^{2t+1-d} \le 2^{2t+1} \sum_{d=0}^{\infty} \binom{d+t-1}{t-1} 2^{-d}$$
$$= 2^{2t+1} \frac{1}{(1-1/2)^t} = 2 \cdot 8^t.$$

Now (9) gives

$$\beta(\Phi_{4,t}) \sim a \cdot b^t, \quad b \approx 14.5682$$

finishing the proof of the theorem. ■

6. Comments and open problems

6.1. Arbitrary fields

We will now indicate how to generalize our construction to an arbitrary field. We first need to review what Brown's hsop looks like over a field F. Fix an orientation of $E(G)$. Also, for each $e_j \in T(G)$, orient all the edges of D_j in one of the two possible directions. Now define signs

$$\epsilon_{i,j} = \begin{cases} 1 & \text{if the orientation of } e_i \text{ in } G \text{ is the same as in } D_j, \\ -1 & \text{if these orientations are opposite.} \end{cases}$$

We have corresponding polynomials

$$\theta_j = \sum_{e_i \in D_j} \epsilon_{i,j} x_i. \tag{15}$$

Theorem 6.1 ([3]). *If G is a connected graph then the set of polynomials defined by (15) for $e \in T(G)$ is an hsop for $F(G)$.* ■

Solving for x_j in the equation for θ_j and plugging into the monomials $\mathbf{x}^{\overline{C}}$, $C \in \mathcal{C}(G)$, give the generators $p_{\overline{C}}$ for an ideal $J(G)$ such that

$$R(G) \cong \frac{F[x_1, \ldots, x_k]}{J(G)}.$$

Note that the monomial $m_{\overline{C}}$ that was chosen from the expansion of $p_{\overline{C}}$ in the case $F = \mathbb{Z}_2$ will also appear with coefficient ± 1 for any field. So the proof of Theorem 3.2 will go through as before as long as the generators of $J(G)$, $J(G\backslash e_q)$, and $J(G/e_q)$ can be related appropriately.

An orientation of G induces orientations of $G\backslash e_q$ and G/e_q merely by keeping each e_i, $i < q$, oriented the same way in all three graphs. Under the assumptions of Theorem 3.2 we showed that $D_j(G\backslash e_q) = D_j(G)\backslash e_k$ for $j > k$. So we can orient $D_j(G\backslash e_q)$ the same way as $D(G)$ in this case. We also have $D_k(G\backslash e_q) = \{e_k\}$, so it does not matter which way we orient e_k in this cut set as x_k is being set to zero in the quotient. Thus we get, as we did in the proof of Theorem 3.2, that the generators for $J(G\backslash e_q)$ are obtained from those for $J(G)$ by setting $x_k = 0$. Similar considerations show that we can define orientations on the cut sets of G/e_q so that the corresponding equalities still hold. So Theorem 3.2 holds, and similarly so do all the rest of the results of the previous sections, over any field.

6.2. Arbitrary graphs

We conjecture that any graph G, with its edge set suitably ordered, has an NBC basis.

Conjecture 6.2. *Let G be a graph. Then there is a standard ordering of $E(G)$ such that $L(G)$ is a basis for $R(G)$.*

We will now outline a possible approach to proving Conjecture 6.2. Even though we have not been able to push it through, it is possible that some of these ideas will be useful in finally proving or disproving this conjecture. Recall that, by (3), it suffices to find a proof when G is a block. But any block other than K_2 (the complete graph on 2 vertices) has a nice recursive structure in that it can be built from a cycle by adding a sequence of paths called *ears*. This result is due to Whitney [17]. Proofs can also be found in the books of Diestel [18, Proposition 3.1.2] and West [19, Theorem 4.2.8].

Theorem 6.3 (*Ear Decomposition Theorem*). *Suppose $G \neq K_2$. Then G is a block if and only if there is a sequence*

$$G_0, G_1, \ldots, G_l = G$$

such that G_0 is a cycle and G_{i+1} is obtained by taking a non-trivial path and identifying its two endpoints with two distinct vertices of G_i. ■

Note that the graph G_1 in the ear decomposition sequence is a theta graph. So one might try to prove Conjecture 6.2 by induction on l, the number of paths added. (Actually, one also needs to induct on the number of edges since one contracts an edge and not a whole path.) In fact, the induction step goes through in much the same way as our proof for theta graphs as long as the path added has length at least two. The difficulty arises if the path is a single edge. In that case, it is still easy to relate the circuits of $G \backslash e_q$, where e_q is the newly added edge, to those of G. But the situation is much more complicated in G/e_q, which may not even be a block. So a more delicate analysis is needed. Unfortunately, there are graphs (such as the complete graphs) where every ear decomposition requires the addition of a single edge at some stage.

6.3. Not quite arbitrary matroids

As a last point, the reader may have noticed that all of the graphical definitions we used to define NBC bases make sense for the broken circuit complex of an arbitrary matroid. So a natural question is whether our construction goes through in that level of generality. Brown, Colbourn, and Wagner [20] have a way of producing an hsop for any representable matroid. (Actually, their construction is of an hsop for the independence complex of the matroid. But this will also give an hsop for the broken circuit complex since it is a subcomplex of the independence complex having the same rank.) So this would be the natural class of matroids in which to look for NBC bases.

6.4. Gröbner bases

We note that, in general, the monomials used to generate $U(G)$ are not the leading terms of a Gröbner basis for the ideal $J(G)$. As an example of this, one can take a theta graph consisting of three paths of length two in the theta labeling as described in Section 4. Then by choosing a suitable orientation for G and its cocircuits, $J(G)$ will have generators

$$\{p_{\overline{C}}\} = \{x_1(x_1+x_2)^2,\ x_2(x_1+x_2)^2,\ x_1x_2^2\} \tag{16}$$

from which we pick monomials

$$\{m_{\overline{C}}\} = \{x_1^3,\ x_2^3,\ x_1x_2^2\} \tag{17}$$

for the NBC basis.

Suppose, towards a contradiction, that there is a term ordering giving (17) as the set of leading terms of a Gröbner basis. Then in that term ordering we either have $x_1 < x_2$ or $x_1 > x_2$. Suppose the former is true. Then x_1^3 is the smallest (monic) polynomial which is homogeneous of degree three.

Also, the generators of $J(G)$ are homogeneous. So if x_1^3 were a leading term of a polynomial in $J(G)$ then, in fact, $x_1^3 \in J(G)$. But it is easy to check that $x_1^3 \notin J(G)$ since it is not a linear combination of the polynomials in (16). (It suffices to consider linear combinations by homogeneity.) One gets a similar contradiction using x_2^3 if one assumes that $x_1 > x_2$. So no such Gröbner basis can exist.

Acknowledgement

We would like to thank David Forge for stimulating discussions.

References

[1] Herbert S. Wilf, Which polynomials are chromatic? in: Colloquio Internazionale sulle Teorie Combinatorie (Roma, 1973), Tomo I, in: Atti dei Convegni Lincei, vol. 17, Accad. Naz. Lincei, Rome, 1976, pp. 247–256.
[2] Hassler Whitney, A logical expansion in mathematics, Bull. Amer. Math. Soc. 38 (1932) 572–579.
[3] Jason I. Brown, Chromatic polynomials and order ideals of monomials, Discrete Math. 189 (1–3) (1998) 43–68.
[4] F.S. Macaulay, Some properties of enumeration in the theory of modular systems, Proc. London Math. Soc. 26 (1927) 531–555.
[5] Richard P. Stanley, The upper bound conjecture and Cohen–Macaulay rings, Studies in Appl. Math. 54 (2) (1975) 135–142.
[6] Ed Swartz, Lower bounds for h-vectors of k-CM, independence, and broken circuit complexes, SIAM J. Discrete Math. 18 (3) (2004) 647–661 (electronic).
[7] Richard P. Stanley, Acyclic orientations of graphs, Discrete Math. 5 (1973) 171–178.
[8] Leslie G. Valiant, The complexity of computing the permanent, Theoret. Comput. Sci. 8((2) (1979) 189–201.
[9] Leslie G. Valiant, The complexity of enumeration and reliability problems, SIAM J. Comput. 8 (3) (1979) 410–421.
[10] Nathan Linial, Hard enumeration problems in geometry and combinatorics, SIAM J. Algebraic Discrete Methods 7 (2) (1986) 331–335.
[11] François Jaeger, Dirk L. Vertigan, Dominic J.A. Welsh, On the computational complexity of the Jones and Tutte polynomials, Math. Proc. Cambridge Philos. Soc. 108 (1) (1990) 35–53.
[12] Wayne Goddard, Claire Kenyon, Valerie King, Leonard J. Schulman, Optimal randomized algorithms for local sorting and set-maxima, SIAM J. Comput. 22 (2) (1993) 272–283.
[13] Ronald L. Graham, Andrew C. Yao, F. Frances Yao, Information bounds are weak in the shortest distance problem, J. Assoc. Comput. Mach. 27 (3) (1980) 428–444.
[14] Udi Manber, Martin Tompa, The effect of number of Hamiltonian paths on the complexity of a vertex-coloring problem, SIAM J. Comput. 13 (1) (1984) 109–115.
[15] Nabil Kahale, Leonard J. Schulman, Bounds on the chromatic polynomial and on the number of acyclic orientations of a graph, Combinatorica 16 (3) (1996) 383–397.
[16] Richard P. Stanley, Enumerative Combinatorics. Vol. 1, in: Cambridge Studies in Advanced Mathematics, vol. 49, Cambridge University Press, Cambridge, 1997, With a foreword by Gian-Carlo Rota, Corrected reprint of the 1986 original.
[17] Hassler Whitney, Congruent graphs and the connectivity of graphs, Amer. J. Math. 54 (1) (1932) 150–168.
[18] Reinhard Diestel, Graph Theory, third ed., in: Graduate Texts in Mathematics, vol. 173, Springer-Verlag, Berlin, 2005.
[19] Douglas B. West, Introduction to Graph Theory, second ed., Prentice Hall Inc., Upper Saddle River, NJ, 1996.
[20] Jason I. Brown, Charles J. Colbourn, David G. Wagner, Cohen–Macaulay rings in network reliability, SIAM J. Discrete Math. 9 (3) (1996) 377–392.