Available online at www.sciencedirect.com
ScienceDirect
Discrete Mathematics 306 (2006) 3253–3266
DISCRETE MATHEMATICS
www.elsevier.com/locate/disc

Graph polynomials from principal pivoting

Roland Glantz, Marcello Pelillo*

Dipartimento di Informatica, Università Ca' Foscari di Venezia, Via Torino 155, 30172 Venezia Mestre, Italy

Received 18 July 2003; received in revised form 29 May 2006; accepted 5 June 2006
Available online 26 July 2006

Abstract

The recursive computation of the *interlace polynomial* introduced by Arratia, Bollobás and Sorkin is defined in terms of a new pivoting operation on undirected simple graphs. In this paper, we interpret the new pivoting operation on graphs in terms of standard pivoting (on matrices). Specifically, we show that, up to swapping vertex labels, Arratia et al.'s pivoting operation on a graph is equivalent to a principal pivot transform on the graph's adjacency matrix, provided that all computations are performed in the Galois field $\mathbb{F}_2$. Principal pivoting on adjacency matrices over $\mathbb{F}_2$ has a natural counterpart on isotropic systems. Thus, our view of the interlace polynomial is closely related to the one by Aigner and van der Holst.

The observations that adjacency matrices of undirected simple graphs are skew-symmetric in $\mathbb{F}_2$ and that principal pivoting preserves skew-symmetry in *all* fields suggest to extend Arratia *et al.*'s pivoting operation to fields other than $\mathbb{F}_2$. Thus, the interlace polynomial extends to polynomials on *gain* graphs, namely bidirected edge-weighted graphs whereby reversed edges carry non-zero weights that differ only by their sign. Extending a proof by Aigner and van der Holst, we show that the extended interlace polynomial can be represented in a non-recursive form analogous to the non-recursive form of the original interlace polynomial, i.e., the Martin polynomial.

For infinite fields it is shown that the extended interlace polynomial does not depend on the (non-zero) gains, as long as they obey a non-singularity condition. These gain graphs are all supported by a single undirected simple graph. Thus, a new graph polynomial is defined for undirected simple graphs. The recursive computation of the new polynomial can be done such that all ends of the recursion correspond to independent sets. Moreover, its degree equals the independence number. However, the new graph polynomial is different from the independence polynomial.

Keywords: Interlace polynomial; Principal pivoting; Independent sets; Independence number; Pivoting polynomial

1. Introduction

Motivated by a problem arising from DNA sequencing by hybridization [2], Arratia, Bollobás and Sorkin introduced the *interlace polynomial*, a graph polynomial defined on undirected graphs (see also [3]). The interlace polynomial can be computed recursively using a new pivoting operation on graphs. The new pivoting operation takes an edge and toggles others (from a non-edge to an edge or vice versa) and, as the authors put it, the situation after just two pivoting operations is already "obscure".

* Corresponding author. Tel.: +39 0412348411; fax: +39 0412348419.
E-mail addresses: glantz@dsi.unive.it (R. Glantz), pelillo@dsi.unive.it (M. Pelillo).

0012-365X/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.disc.2006.06.003

Our motivation for studying pivoting operations on graphs originates from work on graph matching in computer vision and pattern recognition, where the graphs are typically large. The problem of matching two graphs is equivalent to finding large independent sets in an auxiliary graph. The latter problem can be formulated as a linear complementarity problem (LCP), the solution of which involves sequences of pivoting operations on the adjacency matrix of the auxiliary graph [9,8].

The recursive computation of the interlace polynomial in [2] stops whenever an edge-less graph arises, yielding a term x^n in the polynomial, where n is number of vertices in the edge-less graph. One contribution of our paper is a variant of interlace polynomial which can be computed such that any edge of the original graph is preserved. Hence, the edge-less graphs arising from the new polynomial all correspond to independent sets of the original graph. The links between pivoting on graphs and pivoting on adjacency matrices thus establish a new connection between graph matching and graph polynomials. Therefore, an earlier version of this paper has appeared in a volume of the *Lecture Notes in Computer Science* dedicated to graph-based representations in pattern recognition [7].

The three main contributions of the paper are as follows. Firstly, we interpret Arratia et al.'s pivoting operations on graphs in terms of standard pivoting operations on matrices. In particular, we show that pivoting with respect to an edge $e = \{u, v\}$ of an undirected graph G, as defined in [2], is equivalent to a sequence of two classical pivoting steps on an adjacency matrix of G, provided that the computations are performed in the Galois field $\mathbb{F}_2$. If a and b denote the row and column numbers of u and v in an adjacency matrix A of G, respectively, then pivoting w.r.t. e turns out to be equivalent to classical matrix pivoting w.r.t. $A(a, b)$, followed by classical matrix pivoting w.r.t. $A(b, a)$ (here, and in the sequel, $A(a, b)$ denotes the entry of matrix A in row a and column b). This sequence of pivoting operations, in turn, is equivalent to a principal pivot transform on A with respect to the non-singular 2×2 principal submatrix $A[\{a, b\}]$, followed by swapping row a with row b and column a with column b. Here, $A[\{a, b\}]$ denotes the principal submatrix of A given by the rows and columns a, b. Thus, any iteration of pivoting operations on graphs is equivalent to a *single* principal pivoting step on the graph's adjacency matrix, followed by some permutation of the rows and the corresponding columns. Principal pivoting on adjacency matrices over $\mathbb{F}_2$ has a natural counterpart on isotropic systems. Thus, our view of the interlace polynomial is closely related to the one by Aigner and van der Holst. In particular, we will see that the interlace polynomial can be computed in terms of 'principal pivoting' on isotropic systems.

Our second contribution comes from the observation that the principal pivot transform with respect to $A[\{a, b\}]$ preserves skew-symmetry in *all* fields. Note also that A is skew-symmetric in $\mathbb{F}_2$. Hence, by extending the pivoting operation from $\mathbb{F}_2$ to arbitrary fields, we extend the interlace polynomial to gain graphs with gains from an arbitrary field, i.e., to bidirected graphs, the edges of which have attributes g such that $g(r, s) = -g(s, r) \neq 0$ for all edges (r, s). In Arratia et al. [2] the key to prove the existence of the interlace polynomial is a lemma about two consecutive pivoting operations on undirected simple graphs. This lemma turns out to be a special case, i.e., $\mathbb{F} = \mathbb{F}_2$, of an equality easily formulated and proven in terms of principal pivoting on skew-symmetric matrices. Thus, the interlace polynomial readily extends to gain graphs. The original interlace polynomial of an undirected graph G_u equals the Martin polynomial of an isotropic system associated with G_u [1]. The Martin polynomial provides a non-recursive definition of the original interlace polynomial. Extending the work of Aigner et al. to fields other than $\mathbb{F}^2$ we will see that there exists an analogous non-recursive definition for the extended interlace polynomial.

The final part of the paper is devoted to introducing the *pivoting polynomial*, a variant of the interlace polynomial that is again defined on any undirected simple graph G. The plan is to define it via the extended interlace polynomial of gain graphs "corresponding" to G. Here, "corresponding" means that the gain graphs have the same vertices and the same edges (apart from the directions) as G. In addition, the class of gain graphs is restricted by a non-singularity condition on the gains which ensures that the extended interlace polynomial can be computed such that none of the edges in the original gain graph is removed by pivoting. Thus, the vertex sets of the empty graphs arising from selected computations of the extended interlace polynomial are always independent sets of the original gain graphs. It turns out to be crucial that the gains come from an infinite field. We also show that the degree of the pivoting polynomial equals the independence number (whereas the degree of the interlace polynomial is merely an upper bound for it). Finally, we present a new pivoting operation on undirected simple graphs and use it to come up with a recursive form of the pivoting polynomial that involves only undirected simple graphs, pivoting, and taking subgraphs.

The plan of the paper is as follows. Section 2 deals with skew-symmetric matrices only. Here, we study principal and double pivoting on skew-symmetric adjacency matrices over an arbitrary field $\mathbb{F}$. In particular, we express double pivoting in terms of principal pivoting and present the key lemma to prove the existence of the extended interlace polynomial. Here, we also extend pivoting to gain graphs. The first purpose of Section 3 is to show that pivoting on undirected graphs, as defined in [2], corresponds to double pivoting on gain graphs, provided the gains are from $\mathbb{F}_2$. This leads us to extending the interlace polynomial to fields other than $\mathbb{F}_2$ and providing a non-recursive representation of the extended interlace polynomial analogous to the Martin polynomial. In Section 4 we focus on classes of $n \times n$ adjacency matrices of gain graphs, where the gains are from an infinite field $\mathbb{F}$. In particular, we define classes such that: (1) for any pair (M_1, M_2) of matrices from the same class and for any position (i, j), $(1 \leqslant i \leqslant n)$ the entry of M_1 at (i, j) is non-zero, if and only if the entry of M_2 at (i, j) is non-zero, (2) the extended interlace polynomial of two gain graphs is identical whenever they have adjacency matrices in the same class, (3) principal pivoting never turns a non-zero entry in the Schur complement to zero. In Section 5, the pivoting polynomial of an undirected simple graph G_{u} is defined in terms of "corresponding" gain graphs, where "corresponding" is specified via the matrix classes introduced in Section 4. Using a special scheme for calculating the pivoting polynomial it turns out that the ends of the recursive computation of the pivoting polynomial all correspond to independent sets. Here, we also present the new pivoting operation and compare the pivoting polynomial to the original interlace polynomial. Finally, Section 6 deals with deriving the original interlace polynomial via principal pivoting on isotropic systems.

2. Double pivoting on skew-symmetric matrices and gain graphs

Pivoting is a standard method in linear and quadratic optimization [6]. Given a field $\mathbb{F}$ and a matrix $M \in \mathbb{F}^{n\times n}$, the *simple pivot transform* of M w.r.t. an entry $M(a, b) \neq 0$ is the matrix $M^{(a,b)} \in \mathbb{F}^{n\times n}$, where

$$M^{(a,b)}(a, b) = \frac{1}{M(a, b)}, \tag{1}$$

$$M^{(a,b)}(i, b) = \frac{M(i, b)}{M(a, b)}, \quad i \neq a, \tag{2}$$

$$M^{(a,b)}(a, j) = -\frac{M(a, j)}{M(a, b)}, \quad j \neq b, \tag{3}$$

$$M^{(a,b)}(i, j) = \frac{M(i, j) - (M(i, b))}{M(a, b)M(a, j)}, \quad i \neq a, \quad j \neq b. \tag{4}$$

To deal with iterations of simple pivoting operations in a convenient way, we will make use of principal pivot transforms as defined below. In the following, n denotes a positive integer and M is from $\mathbb{F}^{n\times n}$. The next definitions are from [11]:

- $\langle n\rangle := \{1, 2, \ldots, n\}$. For any $\alpha \subseteq \langle n\rangle$ the set $\langle n\rangle \backslash \alpha$ is denoted by $\overline{\alpha}$.
- $M[\alpha, \beta]$ is the submatrix of M whose rows and columns are indexed by α and β, respectively. The submatrix $M[\alpha, \alpha]$ of M is written as $M[\alpha]$. The matrix $M[\alpha]$ is a so-called *principal* submatrix of M.
- The *Schur complement* $M\backslash M[\alpha]$ of a non-singular principal submatrix $M[\alpha]$ in M is defined by

$$M\backslash M[\alpha] := M[\overline{\alpha}] - M[\overline{\alpha}, \alpha]M[\alpha]^{-1}M[\alpha, \overline{\alpha}]. \tag{5}$$

Definition 2.1 (*Principal pivot transform $ppt(M, \alpha)$*). Let $\alpha \subseteq \langle n\rangle$ be such that $M[\alpha]$ is non-singular. The principal pivot transform $ppt(M, \alpha)$ of $M \in \mathbb{F}^{n\times n}$ is obtained from M through replacing

$M[\alpha]$ by $M[\alpha]^{-1}$,
$M[\overline{\alpha}]$ by $M\backslash M[\alpha]$,
$M[\alpha, \overline{\alpha}]$ by $-M[\alpha]^{-1}M[\alpha, \overline{\alpha}]$, and
$M[\overline{\alpha}, \alpha]$ by $M[\overline{\alpha}, \alpha]M[\alpha]^{-1}$.

The following theorems differ from Theorems 3.1 and 3.2 in [11] only in that they are formulated for arbitrary fields $\mathbb{F}$ and not just for the field of complex numbers. However, the proofs given in [11] depend merely on the axioms for fields. The two theorems will serve to derive rules for iterating simple pivot operations.

Theorem 2.2. *Let $M \in \mathbb{F}^{n\times n}$ and $\alpha \subseteq \langle n\rangle$ such that $M[\alpha]$ is non-singular. For each $x, y \in \mathbb{F}^n$ define $u = u(x, y)$, $v = v(x, y) \in \mathbb{F}^n$ by $u_i = y_i$, $v_i = x_i$ for all $i \in \alpha$, and $u_j = x_j$, $v_j = y_j$ for all $j \in \overline{\alpha}$. Then, $N = ppt(M, \alpha)$ is the unique matrix with the property that*

$$Mx = y \quad \leftrightarrow \quad Nu = v. \tag{6}$$

Theorem 2.3. *Let $M \in \mathbb{F}^{n\times n}$ and let $\alpha = \bigcup_{i=1}^{k}\alpha_i \subseteq \langle n\rangle$ for some k such that $\alpha_i \cap \alpha_j = \emptyset$ for all $i \neq j$. If the sequence of matrices*

$$M_0 = M, \quad M_i = ppt(M_{i-1}, \alpha_i), \quad i = 1, \dots, k$$

is well defined, i.e., if all $M_{i-1}[\alpha_i]$ are non-singular, then $ppt(M, \alpha) = M_k$.

From now on we require M to be skew-symmetric, i.e., $M(i, j) = -M(j, i)$ and $M(i, i) = 0$ for all i, j. The former requirement implies the latter whenever $\mathbb{F} \neq \mathbb{F}_2$. From $M(a, b) \neq 0$ it follows $a \neq b$, and

$$M^{(a,b)}(b, j) = M(b, j) \quad \forall j, \tag{7}$$

$$M^{(a,b)}(i, a) = M(i, a) \quad \forall i. \tag{8}$$

In particular, $M^{(a,b)}(b, a) = M(b, a) = -M(a, b) \neq 0$ and thus $(M^{a,b})^{(b,a)}$ is well defined. By symmetry, $(M^{(b,a)})^{(a,b)}$ is also well defined. In the following we write $M^{(a,b)(c,d)}$ instead of $(M^{(a,b)})^{(c,d)}$.

To express $M^{(a,b)(b,a)}$ in terms of principal pivoting, we introduce the following notation. Let $M \in \mathbb{F}^{n\times n}$ and let $1 \leqslant a, b \leqslant n$. Then, M_{ab} denotes the matrix obtained from M by swapping the rows indexed by a and by b, followed by swapping the columns indexed by a and by b.

Proposition 2.4 (*Double and principal pivoting*). *Let $M \in \mathbb{F}^{n\times n}$ be skew-symmetric and let $a \neq b \in \langle n\rangle$ be such that the principal submatrix $M[\{a, b\}]$ is non-singular. Then, $M^{(a,b)(b,a)} = (ppt(M, \{a, b\}))_{ab}$.*

Proof. Eqs. (1)–(8) and Definition 2.1 yield

- $M^{(a,b)(b,a)}(a, a) = 0 = M[\{a, b\}]^{-1}(b, b)$,
  $M^{(a,b)(b,a)}(b, b) = 0 = M[\{a, b\}]^{-1}(a, a)$,
- $M^{(a,b)(b,a)}(a, b) = 1/M(a, b) = M[\{a, b\}]^{-1}(b, a)$,
  $M^{(a,b)(b,a)}(b, a) = 1/M(b, a) = M[\{a, b\}]^{-1}(a, b)$,
- $M^{(a,b)(b,a)}(b, j) = -(M(b, j)/M(b, a)) = (-M[\{a, b\}]^{-1}\, M[\{a, b\}, \overline{\{a, b\}}])(a, j)$ for $j \neq a, b$,
- $M^{(a,b)(b,a)}(a, j) = -(M(a, j)/M(a, b)) = (-M[\{a, b\}]^{-1}M[\{a, b\}, \overline{\{a, b\}}])(b, j)$ for $j \neq a, b$,
- $M^{(a,b)(b,a)}(i, b) = M(i, b)/M(a, b) = (M[\overline{\{a, b\}}, \{a, b\}]M[\{a, b\}]^{-1})(i, a)$ for $i \neq a, b$,
- $M^{(a,b)(b,a)}(i, a) = M(i, a)/M(b, a) = (M[\overline{\{a, b\}}, \{a, b\}]M[\{a, b\}]^{-1})(i, b)$ for $i \neq a, b$,
- $M^{(a,b)(b,a)}(i, j) = M(i, j) - (M(i, b)M(a, j) - M(i, a)M(b, j))/M(a, b)$
  $=(M\backslash M[\{a, b\}])(i, j)$ for $\{i, j\} \cap \{a, b\} = \emptyset$. □

Using Proposition 2.4 and the fact that principal pivoting with respect 2×2-principal submatrices preserves skew-symmetry, we get Proposition 2.5 which, in turn, justifies Definition 2.6.

Proposition 2.5 (*Double pivoting, skew-symmetry*). *Let $M \in \mathbb{F}^{n\times n}$ be skew-symmetric with $M(a, b) \neq 0$. Then, $M^{(a,b)(b,a)} = M^{(b,a)(a,b)}$. Moreover, $M^{(a,b)(b,a)}$ is again skew-symmetric.*

Definition 2.6 (*Double pivot transform $M^{\{a,b\}}$*). Let $M \in \mathbb{F}^{n\times n}$ be skew-symmetric and let $M(a, b) \neq 0$. Then, the skew-symmetric matrix

$$M^{\{a,b\}} := M^{(a,b)(b,a)} \tag{9}$$

is called the double pivot transform of M w.r.t. $\{a, b\}$.

In the following we write $M^{\{a,b\}\{a,c\}}$ instead of $(M^{\{a,b\}})^{\{a,c\}}$.

Lemma 2.7 (*Analogue of Lemma 10(ii) in Arratia et al.* [2]). *Let* $M \in \mathbb{F}^{n\times n}$ *be skew-symmetric with* $M(a,b)$, $M(a,c) \neq 0$. *Then,*

$$M^{\{a,b\}\{a,c\}} = (M^{\{a,c\}})_{bc}. \tag{10}$$

Proof. First we prove that $ppt(ppt(M,\{a,b\}),\{b,c\}) = ppt(M,\{a,c\})$. Applying Theorem 2.2 first to principal pivoting w.r.t. $\{a,b\}$ and then to principal pivoting w.r.t. $\{b,c\}$, it follows that $Mx = y \leftrightarrow Nu = v$, where $N = ppt(ppt(M,\{a,b\}),\{b,c\})$ is unique and u, v are defined by

$$u_i = \begin{cases} y_i & \text{if } i \in \{a,c\}, \\ x_i & \text{otherwise,} \end{cases} \qquad v_i = \begin{cases} x_i & \text{if } i \in \{a,c\}, \\ y_i & \text{otherwise.} \end{cases} \tag{11}$$

Note that x_b and y_b have been exchanged twice. Applying Theorem 2.2 to principal pivoting w.r.t. $\{a,c\}$ only, it follows that $Mx = y \leftrightarrow N'u = v$, where $N' = ppt(M,\{a,c\})$ is unique and u, v are as above. Hence, $N = N'$.

Finally, using Proposition 2.4 and the result just proven, we have

$$\begin{aligned} M^{\{a,b\}\{a,c\}} &= (ppt((ppt(M,\{a,b\}))_{ab},\{a,c\}))_{ac} \\ &= (ppt(ppt(M,\{a,b\}),\{b,c\})_{ab})_{ac} = ((ppt(M,\{a,c\}))_{ab})_{ac} \\ &= (((M^{\{a,c\}})_{ac})_{ab})_{ac} = (M^{\{a,c\}})_{bc}. \qquad \square \end{aligned}$$

In the rest of the paper $G = (V, E, g)$ denotes a gain graph with $V = \{1, \ldots, n\}$ and with non-zero gains. Formally, the edge set E of G is a subset of $(V \times V)\backslash\{(v,v) \mid v \in V\}$ such that $(u,v) \in E$ implies $(v,u) \in E$, and the gains are given by a mapping g from E to $\mathbb{F}\backslash\{0\}$ such that $g(u,v) = -g(v,u)$ for all $(u,v) \in E$. Thus, the adjacency matrix of G (with the entries being the gains if the edge exists and 0 otherwise) is a skew-symmetric matrix, and we can interpret both double and principal pivoting as operations on gain graphs. Specifically, if G has an edge (a,b), and if M is the adjacency matrix of G, the double pivot transform $G^{\{a,b\}} = (V^{\{a,b\}}, E^{\{a,b\}}, g^{\{a,b\}})$ of G is defined by

- $V^{\{a,b\}} := V$,
- $E^{\{a,b\}} := \{(i,j) \mid i, j \in V,\ M^{\{a,b\}}(i,j) \neq 0\}$,
- $g^{\{a,b\}}(i,j) := M^{\{a,b\}}(i,j)$ for all $(i,j) \in E^{\{a,b\}}$.

In particular, Lemma 2.7 extends to

$$G^{\{a,b\}\{a,c\}} = (G^{\{a,c\}})_{bc} \quad \forall (a,b),(a,c) \in E. \tag{12}$$

In accordance with Proposition 2.4 we may also define the principal pivot transform $ppt(G,\{a,b\})$ of G w.r.t. $\{a,b\}$ as

$$ppt(G,\{a,b\}) := (G^{\{a,b\}})_{ab} \quad \forall (r,s) \in E. \tag{13}$$

3. Extending the interlace polynomial to fields other than $\mathbb{F}_2$

In case of $\mathbb{F} = \mathbb{F}_2$, the Galois field containing only the numbers 0 and 1, the simple pivot transform $M^{(a,b)}$ of M w.r.t. $M(a,b) \neq 0$, i.e., w.r.t. $M(a,b) = 1$, takes the form

$$M^{(a,b)}(i,j) = \begin{cases} 1 - M(i,j) & \text{if } i \neq a,\quad j \neq b \text{ and } M(i,b) = M(a,j) = 1, \\ M(i,j) & \text{otherwise.} \end{cases}$$

In Fig. 1 the matrix $M^{(a,b)}$ is interpreted as a directed graph with self-loops. Double pivoting w.r.t. the non-zero entry at $\{a,b\}$ (see again Fig. 1) yields

$$\begin{aligned} M^{\{a,b\}}(i,j) &= M^{(a,b)(b,a)}(i,j) \\ &= \begin{cases} 1 - M^{(a,b)}(i,j) & \text{if } i \neq b,\quad j \neq a \text{ and } M^{(a,b)}(i,a) = M^{(a,b)}(b,j) = 1, \\ M^{(a,b)}(i,j) & \text{otherwise,} \end{cases} \\ &= \begin{cases} 1 - M^{(a,b)}(i,j) & \text{if } i \neq b,\quad j \neq a \text{ and } M(i,a) = M(b,j) = 1, \\ M^{(a,b)}(i,j) & \text{otherwise.} \end{cases} \end{aligned}$$

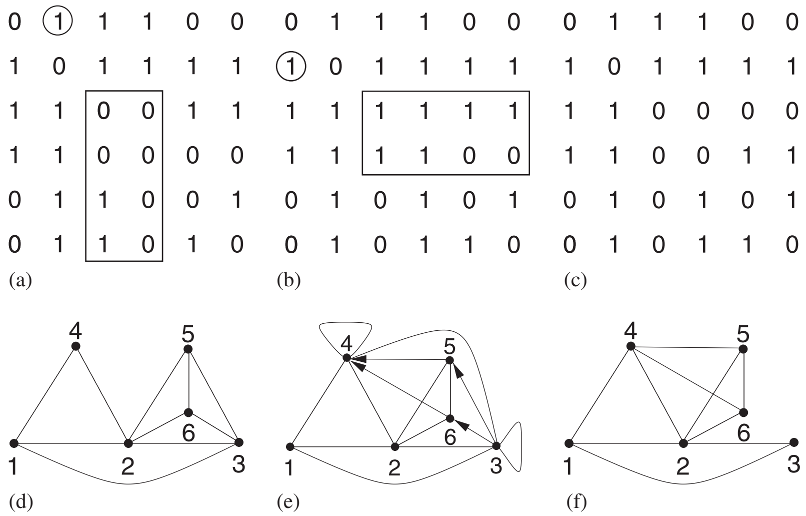

Fig. 1. (a) Adjacency matrix M. The pivoting element (1, 2) is circled. Changes will take place in the rectangle. (b) $M^{(1,2)}$. The new pivoting element is (2, 1). (c) $M^{\{1,2\}}$. (d), (e), (f) Graphs defined by (a), (b), and (c), respectively.

Table 1
Eq. (14) depending on the classes of the nodes i, j

	$j \in C_1^{a,b}$	$j \in C_2^{a,b}$	$j \in C_3^{a,b}$	$j \in C_4^{a,b}$
$i \in C_1^{a,b}$	0 Xor 0	0 Xor 1	0 Xor 1	0 Xor 0
$i \in C_2^{a,b}$	1 Xor 0	0 Xor 0	1 Xor 0	0 Xor 0
$i \in C_3^{a,b}$	1 Xor 0	0 Xor 1	1 Xor 1	0 Xor 0
$i \in C_4^{a,b}$	0 Xor 0	0 Xor 0	0 Xor 0	0 Xor 0

Comparing M with $M^{\{a,b\}}$, the entry at (i, j) changes if and only if the following boolean expression is true:

$$\begin{aligned}&(i \neq a,\ j \neq b,\ M(i,b) = M(a,j) = 1)\\ &\text{Xor}\\ &(i \neq b,\ j \neq a,\ M(i,a) = M(b,j) = 1).\end{aligned} \tag{14}$$

Let $G_u = (V, E_u)$ be an undirected simple graph, i.e., E_u is a collection of two-element subsets of V. G_u corresponds to a bidirected graph $G = (V, E, g)$ with gains in $\mathbb{F}_2$, where $E = \{(i, j) \mid \{i, j\} \in E_u\}$ and $g(e) = 1$ for all $e \in E$. Let (a, b) be an edge of G and let M be the (skew-symmetric) adjacency matrix of G. Partitioning $V \backslash \{a, b\}$ into the classes:

(1) $C_1^{r,s} := \{t \in V \backslash \{r, s\} \mid (t, r) \in E, (t, s) \notin E\}$,
(2) $C_2^{r,s} := \{t \in V \backslash \{r, s\} \mid (t, s) \in E, (t, r) \notin E\}$,
(3) $C_3^{r,s} := \{t \in V \backslash \{r, s\} \mid (t, r), (t, s) \in E\}$,
(4) $C_4^{r,s} := \{t \in V \backslash \{r, s\} \mid (t, r), (t, s) \notin E\}$,

and evaluating the boolean expression (14) according to the membership of i and j to the classes $C_1^{a,b}$ to $C_4^{a,b}$ yields Table 1. Note that $G^{\{a,b\}}$ corresponds to another undirected simple graph. Thus, double pivoting on gain graphs with non-zero gains in $\mathbb{F}_2$ induces an operation on the corresponding undirected simple graphs which may be characterized as follows (see Table 1). An edge $\{i, j\}$ of G_u is toggled (between a non-edge and an edge), if and only if i and j belong to different classes other than class $C_4^{a,b}$. This is precisely how the pivoting operation in [2] is described. Hence we have proven the following.

Proposition 3.1 (*Pivoting operation by Arratia et al.*). *The pivoting operation on undirected graphs used to define the interlace polynomial in* [2] *amounts to double pivoting on the corresponding gain graphs over* $\mathbb{F}_2$.

At this point the question arises as to whether the interlace polynomial defined by Arratia et al. [2] extends to fields other than $\mathbb{F}^2$. The answer will be "yes" (see Theorem 15). Let $\mathbf{G}$ denote the class of gain graphs with $V = \{1, \ldots, n\}$ over a given field, let $G \in \mathbf{G}$, let $G - a$ denote the graph obtained from G by deleting vertex a, and let the order of G be written as $|G|$. Furthermore, the ring of polynomials in a variable x with integer coefficients is denoted by $\mathbf{Z}[\mathbf{x}]$. The following is analogous to [2] (an undirected edge in [2] corresponds to a pair of reversed edges in the gain graph).

Theorem 3.2 (*Extended interlace polynomial*). *There is a unique map* $q : \mathbf{G} \to \mathbf{Z}[\mathbf{x}]$, $G \mapsto q(G)$, *such that*

$$q(G) = \begin{cases} q(G-a) + q(G^{\{a,b\}} - b) & \text{for any edge } (a,b) \text{ of } G, \\ x^{|G|} & \text{if } G \text{ has no edges.} \end{cases} \tag{15}$$

Proof. Using Eq. (12) instead of Lemma 10(ii) in [2], the proof is the same as the in [2]. □

Note that due to Eq. (13) the recursive part of Eq. (15) can be written as

$$q(G) = q(G-a) + q(ppt(G, \{a,b\}) - a). \tag{16}$$

Arratia et al. [4] and Aigner and van der Holst [1] express the interlace polynomial $q(\cdot)$ (the former also a two-variable extension of $q(\cdot)$) in an explicit (non-recursive) way. Specifically, if the vertex set of the (undirected simple) graph G_u is $\{1, \ldots, n\}$, if A_T denotes the adjacency matrix (with entries from $\mathbb{F}_2$) of the subgraph of G_u that is induced by T, and if $\mathrm{co}(A_T)$ denotes the corank of A_T, one can write

$$q(G_u) = \sum_{T \subseteq \{1,\ldots,n\}} (x-1)^{\mathrm{co}(A_T)} \tag{17}$$

(see Eq. (13) in [4] and Eq. (3) in [1]). As noted by Arratia et al. [4] this form of the interlace polynomial suggests an extension to fields other than $\mathbb{F}_2$ (by allowing that the A_T are over a field other than $\mathbb{F}_2$). Now the question arises as to whether the extended interlace polynomial of a gain graph G over $\mathbb{F} \neq \mathbb{F}_2$ in Eq. (15) can also be written as in Eq. (17) (with G instead of G_u). This is indeed the case.

Proposition 3.3 (*Explicit form of* $q(\cdot)$). *Let G be a gain graph with vertex set* $\{1, \ldots, n\}$. *Then the extended interlace polynomial* $q(G)$ *defined recursively in Eq.* (15) *takes the form*

$$q(G) = \sum_{T \subseteq \{1,\ldots,n\}} (x-1)^{\mathrm{co}(A_T)},$$

where A_T *is the skew-symmetric adjacency matrix of the subgraph of G induced by T.*

Proof. Theorem 1 together with Corollary 1 in [1] is a special case, i.e., $\mathbb{F} = \mathbb{F}_2$, of our proposition. The plan is to extend their proof to $\mathbb{F} \neq \mathbb{F}^2$. Consider the $n \times 2n$-matrix $L^G = (A|I)$, i.e., the skew symmetric adjacency matrix of G adjoined to the $n \times n$-identity matrix. If $\mathbb{F} = \mathbb{F}_2$, the rows of L^G form an isotropic system (see [5,1] and Section 6 in this paper). Let the rows of L^G be labeled by $\{1, \ldots, n\}$, and let the columns of L^G be labeled by $\{1, \ldots, n; \bar{1}, \ldots, \bar{n}\}$. As in [1] we call a column set S *admissible* if $|S \cap \{i, \bar{i}\}| = 1$ for all i and let L_S^G denote the $n \times n$-submatrix of L^G with column set S. Due to $\mathrm{co}(A_T) = \mathrm{co}(L_S^G)$ for $T = S \cap \{1, \ldots, n\}$, S admissible (see [1]), we may as well show that

$$q(G) = \sum_S (x-1)^{\mathrm{co}(L_S^G)},$$

where the sum extends over all admissible column sets S. Our proposition now follows from items (1), (2), and (4).

(1) For edge-less G we have $q(G)=x^n$. The proof is identical to the one in [1].

(2) $\sum_{S:\overline{n}\in S}(x-1)^{\mathrm{co}(L_S^G)}=q(G\backslash n)$. The proof is identical to the one in [1].

(3) Let S' denote the (admissible) set obtained from S by swapping the membership of $n-1$ and $\overline{n-1}$, as well as that of n and $\overline{n}$. Then, $\mathrm{co}(L_S^G)=\mathrm{co}(L_{S'}^{ppt(G,\{n-1,n\})})$ for all admissible S. Indeed, extending the proof in [1] we may write

$$L^G=\left(\begin{array}{ccc|ccc} B & C & D & I & 0 & 0\\ -C^T & 0 & a & 0 & 1 & 0\\ -D^T & -a & 0 & 0 & 0 & 1\end{array}\right),$$

where B is a skew-symmetric $n-2\times n-2$-matrix, I is the $n-2\times n-2$-identity matrix, C, D are column vectors with $n-2$ entries, and a is a single non-zero entry. Similar to the proof in [1] we multiply L^G from the left by the invertible $n\times n$-matrix

$$M=\begin{pmatrix} I & -\frac{1}{a}D & \frac{1}{a}C\\ 0 & \frac{1}{a} & 0\\ 0 & 0 & -\frac{1}{a}\end{pmatrix}$$

and get

$$ML^G=\left(\begin{array}{ccc|ccc} B+\frac{1}{a}DC^T-\frac{1}{a}CD^T & 0 & 0 & I & -\frac{1}{a}D & \frac{1}{a}C\\ -\frac{1}{a}C^T & 0 & 1 & 0 & \frac{1}{a} & 0\\ \frac{1}{a}D^T & 1 & 0 & 0 & 0 & -\frac{1}{a}\end{array}\right).$$

As in [1] we swap columns $n-1$, $\overline{n-1}$, columns n, $\overline{n}$, and rows $n-1$, n. Finally, we multiply the columns $n-1, n, \overline{n-1}, \overline{n}$ and the rows $n-1, n$ by -1. We end up with the matrix

$$\left(\begin{array}{ccc|ccc} B+\frac{1}{a}DC^T-\frac{1}{a}CD^T & \frac{1}{a}D & -\frac{1}{a}C & I & 0 & 0\\ -\frac{1}{a}D^T & 0 & -\frac{1}{a} & 0 & 1 & 0\\ \frac{1}{a}C^T & \frac{1}{a} & 0 & 0 & 0 & 1\end{array}\right)=L^{ppt(G,\{n-1,n\})}. \tag{18}$$

From M being non-singular it now follows $\mathrm{co}(L_S^G)=\mathrm{co}(ML_S^G)$ and thus $\mathrm{co}(L_S^G)=\mathrm{co}(L_{S'}^{ppt(G,\{n-1,n\})})$ for all admissible S.

(4) $\sum_{S:n\in S}(x-1)^{\mathrm{co}(L_S^G)}=q(G^{\{n-1,n\}}\backslash n-1)$. Indeed, using Eq. (13), we may as well prove that

$$\sum_{S:n\in S}(x-1)^{\mathrm{co}(L_S^G)}=q(ppt(G,\{n-1,n\})\backslash n). \tag{19}$$

Looking at the matrix in Eq. (18) one sees that the $\overline{n}$th column of $L^{ppt(G,\{n-1,n\})}$ has only one non-zero entry, i.e., the entry 1 in the last row. Thus, for any admissible S' with $n\in S'$, the last row of $L_{S'}^{ppt(G,\{n-1,n\})}$ (equal to the nth row of $L^{ppt(G,\{n-1,n\})}$) cannot be written as a linear combination of the rows $1,\ldots,n-1$. We conclude that $rank(L_{S'}^{ppt(G,\{n-1,n\})})=1+rank(L_{S'}^{ppt(G,\{n-1,n\})\backslash n})$, i.e., that $\mathrm{co}(L_{S'}^{ppt(G,\{n-1,n\})})=\mathrm{co}(L_{S'}^{ppt(G,\{n-1,n\})\backslash n})$. From Item (3) above it now follows that $\mathrm{co}(L_{S'}^{ppt(G,\{n-1,n\})\backslash n})=\mathrm{co}(L_S^G)$ for all admissible S. By induction we may assume that Eq. (19) has already been shown for all gain graphs with less than n vertices. This yields Eq. (19). $\square$

4. Classes of skew-symmetric adjacency matrices

The plan for the rest of the paper is to derive a new graph polynomial for undirected simple graphs $G_u = (V, E_u)$ via the extended interlace polynomial of "corresponding" gain graphs. The term "corresponding" implies that the gain graphs take the form $G = (V, E, g(\cdot))$, where $E = \{(u, v) \mid \{u, v\} \in E_u\}$ and $g(e) \neq 0$ for all $e \in E$. We will have to make sure that from G and G' both corresponding to G_u, it always follows that $q(G) = q(G')$. The new graph polynomial will be different from the interlace polynomial defined by Arratia et al. [2].

In this section the problem of finding "corresponding" gain graphs is approached via classes of skew-symmetric adjacency matrices. To formalize our approach we proceed as follows.

Definition 4.1 (*Supports $\underline{M}$, $\underline{\mathbf{M}}$*). Let $M \in \mathbb{F}^{n\times n}$. The support $\underline{M} \in \{0, 1\}^{n\times n}$ of M is defined by

$$\underline{M}(i, j) := \begin{cases} 1 & \text{if } M(i, j) \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\mathbf{M}$ be a collection of matrices from $\{0, 1\}^{n\times n}$. Then, the support of $\mathbf{M}$ is the $n \times n$ matrix

$$\underline{\mathbf{M}}(i, j) := \begin{cases} 1 & \text{if there exists } M \in \mathbf{M} \text{ with } M(i, j) = 1, \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

Given an adjacency matrix $A \in \{0, 1\}^{n\times n}$ of G_u, we are looking for a class $\max^A$ of skew-symmetric adjacency matrices from $\mathbb{F}^{n\times n}$ such that

(1) $\underline{M} = A$ for all $M \in \max^A$.
(2) Let $M, M' \in \max^A$ be skew-symmetric adjacency matrices of the gain graphs G and G'. Then, $q(G) = q(G')$.

Provided that $\mathbb{F}$ is infinite, the non-empty class $\max^A$ defined at the end of this section will turn out to solve the problem. Moreover, we will see that principal pivoting on a matrix from $\max^A$ never turns a non-zero entry of the Schur complement to zero. As shown in the next section the latter result implies that the recursive computation of $q(G)$ can be done such that each end of the recursion corresponds to an independent set of G.

For the time being we restrict ourselves to principal pivoting on skew-symmetric $n \times n$ adjacency matrices with a fixed support A and with respect to a fixed $\alpha \subseteq \langle n\rangle$. Of course, α has to be such that there exists a skew-symmetric matrix M with non-singular $M[\alpha]$ and $\underline{M} = A$.

Definition 4.2 (*Regular w.r.t. A, n_A*). Let $A \in \{0, 1\}^{n\times n}$ be symmetric with $A(i, i) = 0$ for all i. The set $\alpha \subseteq \langle n\rangle$ is said to be regular w.r.t. A, if there exists skew-symmetric $M \in \mathbb{F}^{n\times n}$ such that $\underline{M} = A$ and $M[\alpha]$ is non-singular. Furthermore, set

$$n_A := \{\alpha \subseteq \langle n\rangle \mid \alpha \text{ is regular with respect to } A\}.$$

We assume that α is regular with respect to A, set

$$\mathbf{V}_{A,\alpha} := \{N \in \mathbb{F}^{n\times n} \mid N \text{ is skew-symmetric, } N[\alpha] \text{ is non-singular, } \underline{N} = A\}$$

and ask whether the support of

$$\mathbf{T}_{A,\alpha} := \{\underline{N\backslash N[\alpha]} \mid N \in \mathbf{V}_{A,\alpha}\} \tag{21}$$

is contained in $\mathbf{T}_{A,\alpha}$. The answer will be "yes, provided that $\mathbb{F}$ is infinite". In this case the elements in $\mathbf{V}_{M,\alpha}$ that yield elements of $\mathbf{T}_{A,\alpha}$ whose support is not the support of $\mathbf{T}_{A,\alpha}$ turn out to be singular cases. In other words, the regular case is the one with a maximum number of 1-entries in $\mathbf{T}_{A,\alpha}$.

Proposition 4.3 ($\mathbf{T}_{A,\alpha} \in \mathbf{T}_{A,\alpha}$). *Let $\mathbb{F}$ be infinite, let $A \in \{0, 1\}^{n\times n}$ be symmetric with $A(i, i) = 0$ for all i, and let $\alpha \in n_A$. Then, $\underline{\mathbf{T}_{A,\alpha}} \in \mathbf{T}_{A,\alpha}$ and $\underline{\mathbf{T}_{A,\alpha}}(i, j) = 0$ implies $A(i, j) = 0$ for all $i, j \in \overline{\alpha}$.*

Proof. First we prove that for all $N, N' \in \mathbf{V}_{A,\alpha}$ such that $N\backslash N[\alpha](i,j)=0$ and $N'\backslash N'[\alpha](i,j) \neq 0$ for some $i, j \in \overline{\alpha}$, there exists $N'' \in \mathbf{V}_{A,\alpha}$ such that $N''\backslash N''[\alpha](i,j) \neq 0$, and $N''\backslash N''[\alpha](i,j)=0$ implies $N\backslash N[\alpha](i,j)=0$. Indeed, from

$$\begin{aligned} 0 &\neq (N'\backslash N'[\alpha])(i,j) \\ &= N'(i,j) - (N'[\overline{\alpha},\alpha]N'[\alpha]^{-1}N'[\alpha,\overline{\alpha}])(i,j) \\ &= N'(i,j) - \sum_{k\in\alpha}\left(N'(i,k)\left(\sum_{m\in\alpha} N'[\alpha]^{-1}(k,m)N'(m,j)\right)\right) \end{aligned}$$

it follows that $N'(i,j) \neq 0$ or that there exist $k, m \in \alpha$ such that $N'(i,k)$, $N'[\alpha]^{-1}(k,m)$, and $N'(m,j)$ are non-zero. In the first case all but finitely many coupled modifications of $N(i,j)$ and $N(j,i)$ (the resulting matrix has to be skew-symmetric again) yield matrices N'' that fulfill the conditions above. In the second case all but finitely many coupled modifications of $N(i,k)$ and $N(k,i)$ yield matrices N'' that fulfill the conditions above (note that $N(k,i)$ has no effect on the value of $N(i,j)$, since otherwise $i=j$, a contradiction to $N'(i,j)\neq 0$). The first part of the proof now follows from $\mathbb{F}$ being infinite.

Looking at the first part of the proof again, we see that $N''(i,j)\neq 0$ implies $N''\backslash N''[\alpha](i,j)\neq 0$.

Repeatedly applying the modifications in the first part of the proof we can construct $N^* \in \mathbf{V}_{M,\alpha}$ such that $\mathbf{T}_{M,\alpha} \ni \underline{N^*\backslash N^*[\alpha]} = \underline{\mathbf{T}_{A,\alpha}}$. □

From now on the field $\mathbb{F}$ is always infinite. Proposition 4.3 implies that the matrix class defined below is non-empty.

Definition 4.4 (*Matrix class* $\max_{A,\alpha}$). Let $A \in \{0,1\}^{n\times n}$ be symmetric with $A(i,i)=0$ for all i and let $\alpha \in N_A$. Then, $\max_{A,\alpha}$ consists of all skew-symmetric $N \in \mathbb{F}^{n\times n}$ such that $\underline{N}=A$, $N[\alpha]$ is non-singular, and $\underline{N\backslash N[\alpha]} = \underline{\mathbf{T}_{N,\alpha}}$.

Finally, the matrix class $\max^A$ is defined as follows.

Definition 4.5 (*Matrix class* $\max^A$). Let $A \in \{0,1\}^{n\times n}$ be symmetric with $A(i,i)=0$ for all i. Then, the class max^A is defined by

$$\max^A := \bigcap_{\alpha\in n_A} \max_{A,\alpha}.$$

Looking at Proposition 4.3 once again, we get the following.

Proposition 4.6. *Let $M, N \in \max^A$ and let $\alpha \in n_A$. Then, $\underline{(ppt(M,\alpha))(i,j)} = \underline{(ppt(N,\alpha))(i,j)}$ and $(ppt(M,\alpha))(i,j)=0$ implies $M(i,j)=0$ for all $i, j \in \overline{\alpha}$.*

5. The pivoting polynomial

The following definition allows us to extend the results of Section 4 to classes of gain graphs.

Definition 5.1 (*Maximum gain graphs corresponding to* G_u). Let G_u be an undirected simple graph with vertex set V. Moreover, let A be an adjacency matrix of G_u. A gain graph with vertex set V is called maximum gain graph corresponding to G_u, if it has an adjacency matrix in $\max^A$.

In the following we will see that the extended interlace polynomials of maximum gain graphs corresponding to G_u are identical.

A gain graph G' is called *first-order-descendant* of G, if G has an edge (a,b) such that $G'=G-a$ or $G'=G^{\{a,b\}}-b$. A *descendant of order* $n+1$ is a first-order-descendant of an nth order descendant. It will turn out to be convenient if G may be referred to as its descendant of order 0. G' is called a *descendant* of G, if G' is an nth order descendant of G for some $n \in \mathbb{N}_0$. A descendant of G is called *terminal*, if it has no edges (and thus no descendants).

Proposition 5.2 (*Descendants of $q(G)$*). *Let G be a gain graph. Then, the recursive computation of $q(G)$ can be done such that each descendant is a subgraph of G or a subgraph of some $G^{U_1,\ldots,U_m}\backslash U$, where $U=\bigcup_{i=1}^{m} U_i \subseteq V$ and each U_i has cardinality* 2.

Proof. Consider the following scheme for computing $q(G)$.

- *Step $i=1$*. Pick a non-isolated vertex a_1 of G. As long as a_1 is a non-isolated vertex of G, continue the computation of $q(G)$ via the descendants $G-a_1$ and $G^{\{a_1,b_1\}}-b_1$ for some b_1. The descendant $G-a_1$ has no further descendants at this stage of the scheme. If a_1 is not isolated in $G^{\{a_1,b_1\}}-b_1$, the scheme continues with the computation of $G^{\{a_1,b_1\}}-b_1$ by selecting the descendants $G^{\{a_1,b_1\}}-a_1-b_1$ and $(G^{\{a_1,b_1\}}-b_1)^{\{a_1,b_2\}}-b_2$ for some b_2. Again, the first descendant has no further descendants at this stage and, using Eq. (12), the second descendant can be written as $(G^{\{a_1,b_1\}}-b_1)^{\{a_1,b_2\}}-b_2=$

  $$(G^{\{a_1,b_1\}})^{\{a_1,b_2\}}-b_1-b_2=G^{\{a_1,b_2\}}-b_1-b_2.$$

  In the following let d denote the degree of a_1 in G. Continued expansion of each second descendant eventually leaves us with the descendant $G-a_1$, descendants of the form

  $$G^{\{a_1,b_k\}}-a_1-b_1-b_2-\cdots-b_k \quad (k<d), \tag{22}$$

  and the descendant $G^{\{a_1,b_d\}}-b_1-\cdots-b_d$ with isolated a_1. The expression of the latter descendant can be simplified as follows. Pivoting w.r.t. $\{a_1,b_d\}$ has no effect on the neighborhood of a_1. Hence, $b_1,\ldots,b_d$ are precisely the neighbors of a_1 in G and b_d is the only neighbor of a_1 in $G-b_1-\cdots-b_{d-1}$. Thus, $G^{\{a_1,b_d\}}-b_1-\cdots-b_d=(G-b_1-\cdots-b_{d-1})^{\{a_1,b_d\}}-b_d=$

  $$G-b_1-\cdots-b_d.$$

  Hence, each descendant is a subset of G or it has form (22).
- *Steps $i>1$*. We may assume that each descendant D is a subgraph of G or a subgraph of $G^{U_1,\ldots,U_j}-\bigcup_{k=1}^{j} U_k$ for some U_k, $1\leqslant k\leqslant j\leqslant i-1$. If D has a non-isolated vertex a_i, we apply the first step of the scheme to a_i in D instead of a_1 in G. All descendants still have the required form. □

Theorem 5.3 (*Identical polynomials, independent sets*). *Let G and G' be maximum gain graphs corresponding to G_{u}. Then $q(G)=q(G')$. Furthermore, for each term $a_k x^k$ of $q(G)$ with $a_k\neq 0$ there exist independent sets of G_{u} with cardinality k.*

Proof. According to Proposition 5.2 the polynomial $q(G)$ can be computed such that every descendant is a subgraph of G or a subgraph of some $G^{U_1,\ldots,U_m}\backslash U$, where $U=\bigcup_{i=1}^{m} U_i$, the U_i are disjoint, and each U_i has cardinality 2.

If A is an adjacency matrix of G_{u}, then G has an adjacency matrix $M\in\max^A$. From Theorem 2.3 it follows that $G^{U_1,\ldots,U_m}\backslash U$ has an adjacency matrix equal to $M\backslash M[\alpha]$ for some $\alpha\subseteq\langle|G|\rangle$. Since $\underline{N\backslash N[\alpha]}$ is unique for $N\in\max^A$ (see Proposition 4.6), it follows that $q(G)$ and $q(G')$ can be computed according to the same scheme (of the type specified in the proof of Proposition 5.2) and that the supports of the descendants are the same for G and G'. Hence $q(G)=q(G')$.

Moreover, Proposition 4.6 implies that the edge set of any $G^{U_1,\ldots,U_m}\backslash U$ is a superset of the edge set of $G\backslash U$. Hence, computing $q(G)$ according to the scheme in the proof of Proposition 5.2, the (edge-less) terminal descendants of $q(G)$ correspond to independent sets of G_{u}. □

Proposition 5.3 implies that the following is well defined.

Definition 5.4 (*Pivoting polynomial $p(G_{\mathrm{u}})$*). Let G_{u} be an undirected simple graph and let G be a maximum gain graph corresponding to G_{u}. The pivoting polynomial $p(G_{\mathrm{u}})$ of G_{u} is defined by $p(G_{\mathrm{u}}):=q(G)$.

The following is a consideration from [2] (extended from undirected simple graphs to gain graphs). Consider a vertex a_1 of a gain graph G. If a_1 is an isolated vertex of G, then $q(G)=xq(G-a_1)$. Otherwise, for any edge (a_1,b) it holds that $q(G) = q(G - a_1) + q(G^{\{a_1,b\}} - b)$. In any case, since all coefficients of all $q(\cdot)$ are non-negative, it follows that the degree of $q(G)$ is not smaller than the degree of $q(G - a_1)$. By subsequently taking a_i that are not contained in a fixed maximum independent set, we get that the degree of $q(G)$ is an upper bound for the independence number.

From Proposition 5.3 it follows that the degree of $q(G)$ is a lower bound for the independence number, if G is a maximum gain graph corresponding to some G_{u}. Hence, we get the following.

Theorem 5.5 (*Degree of $p(\cdot)$ equals the independence number*). *For any undirected, simple graph G_{u} the degree of $p(G_{\mathrm{u}})$ is equal to the independence number of G_{u}.*

The example $p(K^3) = 4x$ shows that the pivoting polynomial, in contrast to the independence polynomial [10], may list an independent set more than once.

Recall that the original recursive definition of Arratia et al.'s interlace polynomial is

$$q(G_{\mathrm{u}}) := q(G_{\mathrm{u}} - a) + q(G_{\mathrm{u}}^{\{a,b\}} - b),$$

where a, b is an edge of G_{u} and $G_{\mathrm{u}}^{\{a,b\}}$ is another undirected simple graph that one can derive easily from G_{u}. Now the question arises as to whether $p(G_{\mathrm{u}})$ can be computed in a similar way, i.e., without resorting to a maximum gain graph corresponding to G_{u}, or equivalently, to a skew-symmetric matrix $M \in \max^A$, where $A \in \{0, 1\}^{n \times n}$ is the adjacency of G_{u}.

Fortunately, it is possible to characterize the supports of the descendants of M without knowing M (i.e., just knowing that $M \in \max^A$). Indeed, let $M \in \max^A$ and assume that $M(a, b) \neq 0$. Then, according to the last item in the proof of Proposition 2.4, a zero-entry $M(i, j)$, $i, j \in \bar{\alpha}$, turns into a non-zero-entry $M^{\{a,b\}}(i, j)$ whenever

$$(M(a, i) \neq 0 \wedge M(b, j) \neq 0) \vee (M(a, j) \neq 0 \wedge M(b, i) \neq 0),$$

and these are the only changes between M and $M^{\{a,b\}}$. This leads us to a new pivoting operation on G_{u} and a recursive formula for $q(\cdot)$ in terms of the new operation.

Definition 5.6. Let $G_{\mathrm{u}} = (V, E)$ be an undirected graph and let $\{a, b\} \in E$. Then, $G_{\mathrm{u}}^{[a,b]} := (V^{[a,b]}, E^{[a,b]})$ is given by $V^{[a,b]} = V$ and

$$E^{[a,b]} = E \cup \{\{i, j\} \mid \{a, i\}, \{b, j\} \in E \text{ or } \{a, j\}, \{b, i\} \in E\}.$$

Proposition 5.7 ($p(G_{\mathrm{u}})$ *in terms of* $G_{\mathrm{u}}^{\{a,b\}}$). *Let $\{a, b\}$ an edge of G_{u}. Then*

$$p(G_{\mathrm{u}}) = p(G_{\mathrm{u}} - a) + p(G_{\mathrm{u}}^{[a,b]} - b).$$

It is clear that up to three vertices in G_{u} we have $p(G_{\mathrm{u}}) = q(G_{\mathrm{u}})$. There exists a graph G_4 with four vertices, e.g., the graph with adjacency matrix

$$A = \begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}$$

such that $p(G_4) \neq q(G_4)$. In particular, $p(G_4) = 6x + x^2$ while $q(G_4) = 4x + 2x^2$.

For the complete graph K^n we have $(K^n)^{[a,b]} = (K^n)^{\{a,b\}} = K^n$. Thus, $p(K^n) = q(K^n)$. Likewise, let P_n be the path with n vertices. Then, $(P_n)^{[a,b]} = (P_n)^{\{a,b\}}$ and $p(P_n) = q(P_n)$. For explicit formulas see [2].

6. Principal pivoting and isotropic systems

In [1] the interlace polynomial as introduced by Arratia et al. [2], i.e., $q(G)$ from Theorem 3.2 when $\mathbb{F} = \mathbb{F}_2$, is shown to coincide with the Martin polynomial $m(\mathscr{S}_G)$. Specifically,

$$\mathscr{S}_G = (V, \mathscr{L}_G)$$

is an isotropic system (see below), where $V=\{1,\ldots,n\}$ is the vertex set of G, and $\mathscr{L}_G$ is the row space of the adjacency matrix M of G adjoined to the $n\times n$ identity matrix I_n. Hence, $\mathscr{L}_G$ is the row space of the $n\times 2n$ matrix

$$L_G := (M|I_n).$$

For the following we need the bilinear form $\langle\cdot,\cdot\rangle$ on $\mathbb{F}_2^2=\{(0,0),(0,1),(1,0),(1,1)\}$ defined by

$$\langle (x,y),(x',y')\rangle := \begin{cases} 1 & \text{if } (0,0)\neq(x,y)\neq(x',y')\neq(0,0),\\ 0 & \text{otherwise.}\end{cases}$$

This bilinear form extends to $\mathbb{F}_2^{2n}$ by

$$\langle R,R'\rangle = \sum_{v\in V}\langle (R_v,R_{n+v}),(R_v,R_{n+v})\rangle.$$

Aigner and van der Holst [1] show that $\mathscr{S}_G=(V,\mathscr{L}_G)$ is an isotropic system (see also [5]). This means that

- $\dim(\mathscr{L}_G)=n$ and
- $\langle R,R'\rangle=0$ for all rows R, R' of L_G.

Following Aigner and van der Holst [1] the Martin polynomial of $\mathscr{S}_G$ is given by

$$m(\mathscr{S}_G)=\sum_C (x-1)^{\dim(\mathscr{L}_G\cap\mathscr{C})},$$

where the sum is extended over all rows C with $(C_v,C_{v+n})\neq(0,0),(1,1)$, and $\mathscr{C}$ denotes the set of all restrictions of C. Specifically, R is a restriction of C, if there exists $P\subseteq V$ such that

$$(R_v,R_{v+n}) = \begin{cases} (C_v,C_{v+n}) & \text{if } v\in P,\\ (0,0) & \text{if } v\notin P.\end{cases}$$

Since G has no self-loops, the rows of L_G are complementary, i.e., $R(v)=1$ implies $R(v+n)=0$ for all v and for all rows R. Let the kth column of I_n be denoted by u_k. Then, Mu_k coincides with the kth column of M and, due to M being symmetric, also with the kth row of M. In other words, the complementarity of the rows can also be seen as complementarity between u_k and Mu_k, i.e., as in a solution of a LCP (see [6]).

Theorem 2.2 implies that principal pivoting preserves complementarity in the following sense. Let $x\in\mathbb{F}_2^n$ and $y:=Mx$ be complementary, let $M[\alpha]$ be non-singular, and let

$$u_i=\begin{cases} y_i & \text{if } i\in\alpha,\\ x_i & \text{otherwise.}\end{cases}\qquad v_i=\begin{cases} x_i & \text{if } i\in\alpha,\\ y_i & \text{otherwise.}\end{cases}$$

Then $ppt(M,\alpha)u=v$ and u, v are complementary. Proposition 5.2 states that $q(G)$ can be computed via descendants whose adjacency matrices all take the form $ppt(M,\alpha)[\overline{\alpha}]$ for some α such that $M[\alpha]$ is non-singular.

We now consider the $n\times 2n$ matrix L_D for one such descendant D. It takes the form $ppt(M,\alpha)[\overline{\alpha}]|I_m$, where $m=|\overline{\alpha}|$ is the number of vertices in D. Thus, the isotropic system $\mathscr{S}_D$ can be obtained from the isotropic system $\mathscr{S}_G$ by a *single* principal pivoting step w.r.t. some α, followed by a deletion of all α-rows and α-columns.

7. Conclusions

By expressing pivoting on graphs in terms of principal pivoting on adjacency matrices, we have been able to extend the interlace polynomial from undirected simple graphs to gain graphs over arbitrary fields. Provided that the underlying field is infinite, gain graphs with identical vertex and edge sets were shown to yield the same extended interlace polynomial, as long as they fulfill a non-singularity condition formulated in terms of their (non-zero) gains. Thus, the pivoting polynomial, a new graph polynomial on undirected simple graphs could be defined via gain graphs that fulfill the non-singularity condition. The recursive computation of the pivoting polynomial can be done such that each end of the recursion corresponds to an independent set, one of which is maximum. In contrast to the independence polynomial the pivoting polynomial may list an independent set more than once.

Acknowledgments

We would like to thank the reviewer who pointed out the connection between principal pivoting and isotropic systems as outlined in Section 6. The same reviewer also pointed out that a generalization of the interlace polynomial from $\mathbb{F}_2$ to other fields had already been suggested by Arratia et al. [4] (see Section 3).

References

[1] M. Aigner, H. van der Holst, Interlace polynomials, Linear Algebra Appl. 377 (2004) 11–30.
[2] R. Arratia, B. Bollobás, G.B. Sorkin, The interlace polynomial: a new graph polynomial, Technical Report RC 21813 (98165), IBM T.J. Watson Research Center, Yorktown Heights, NY, July 2000.
[3] R. Arratia, B. Bollobás, G.B. Sorkin, The interlace polynomial: a new graph polynomial, in: Proceeding of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms, San Francisco, CA, January 2000, pp. 237–245.
[4] R. Arratia, B. Bollobás, G.B. Sorkin, A two–variable interlace polynomial, 2002.
[5] A. Bouchet, Tutte–Martin polynomials and orienting vectors of isotropic systems, Graphs Combin. 7 (1991) 235–252.
[6] R.W. Cottle, J.-S. Pang, R.E. Stone, The Linear Complementarity Problem, Academic Press, Boston, MA, 1992.
[7] R. Glantz, M. Pelillo, Graph polynomials, principal pivoting, and maximum independent sets, in: E.R. Hancock, M. Vento (Eds.), Graph Based Representations in Pattern Recognition, Fourth IAPR International Workshop, GbRPR 2003, York, UK, June 30–July 2, 2003, Proceedings, Lecture Notes in Computer Science, vol. 2726, Springer, Berlin, 2003, pp. 166–177.
[8] A. Massaro, M. Pelillo, Matching graphs by pivoting, Pattern Recognition Lett. 24 (2003) 1099–1106.
[9] A. Massaro, M. Pelillo, I.M. Bomze, A complementary pivoting approach to the maximum weight clique problem, SIAM J. Optim. 12 (4) (2002) 928–948.
[10] R.P. Stanley, Enumerative Combinatorics, vol. 1, Wadsworth and Brooks-Cole, Monterey, CA, 1986.
[11] M.J. Tsatsomeros, Principal pivot transforms: properties and applications, Linear Algebra Appl. 307 (2000) 151–165.