Available online at www.sciencedirect.com
SCIENCE DIRECT
ACADEMIC PRESS
Journal of Combinatorial Theory, Series B 88 (2003) 377–398
Journal of Combinatorial Theory Series B
http://www.elsevier.com/locate/jctb

The linear delta-matroid parity problem

James F. Geelen,[a,1] Satoru Iwata,[b,2] and Kazuo Murota[b,2]

[a] Department of Combinatorics and Optimization, University of Waterloo, Waterloo N2L 3G1, Canada
[b] Department of Mathematical Informatics, Graduate School of Information Science and Technology, University of Tokyo, Tokyo 113-8656, Japan

Received 8 July 1997

Abstract

This paper addresses a generalization of the matroid parity problem to delta-matroids. We give a minimax relation, as well as an efficient algorithm, for linearly represented delta-matroids. These are natural extensions of the minimax theorem of Lovász and the augmenting path algorithm of Gabow and Stallmann for the linear matroid parity problem.

1. Introduction

The matroid parity problem is an extremely general problem; containing matching, matroid intersection, as well as NP-hard problems. It is even known to have exponential lower bounds in an ordinary oracle model [11,13]. We consider the parity problem in the more general setting of delta-matroids.

Bouchet [1] introduced delta-matroids as a generalization of matroids [18,20]. Essentially equivalent combinatorial structures were proposed independently by Dress and Havel [7] and by Chandrasekaran and Kabadi [6]. This generalization maintains nice matroidal properties with respect to linear optimization; such as the greedy algorithm and the polyhedral description [4].

In contrast to the case of matroids, where the intersection theorem due to Edmonds [8] has played an important rôle, there are negative results on a pair of delta-matroids. For instance, Chandrasekaran and Kabadi [6] remarked that the

[1] Part of this research was done while this author visited Research Institute for Mathematical Sciences, Kyoto University, supported by Japan Society for Promotion of Science.

[2] This research was done while this author was at Research Institute for Mathematical Sciences, Kyoto University.

0095-8956/03/$ - see front matter
doi:10.1016/S0095-8956(03)00039-X

intersection problem for delta-matroids includes the matroid parity problem, which cannot be solved in polynomial time [13]. Alternatively, Bouchet [3] introduced the delta-covering problem, which is again a generalization of the matroid parity problem and hence is polynomially unsolvable in general.

The matroid parity problem is solvable if the matroid in question is linearly represented. This is a result of Lovász [12–14], who discovered a minimax relation as well as a polynomial-time algorithm. Since then, more efficient algorithms have been developed for the linear matroid parity problem. Gabow and Stallmann [9] devised an augmenting path algorithm, which provides an alternative proof of the minimax theorem of Lovász. In addition, Orlin and Vande Vate [17] presented another efficient algorithm, which relies on the minimax relation. These results motivate us to focus on the delta-covering problem for linear delta-matroids (these are delta-matroids that are represented by skew-symmetric matrices).

In this paper, we introduce the delta-parity problem, which is equivalent to the delta-covering problem but generalizes the matroid parity problem more directly. We then extend the Lovász minimax theorem to the delta-parity problem for linear delta-matroids. The proof follows from an analysis of an augmenting path algorithm, which extends the linear matroid parity algorithm of Gabow and Stallmann. It is not difficult to convert the minimax theorem and the augmenting path algorithm to those for the delta-covering problem. This resolves open problems of Bouchet [3, Problems 14–16] on special classes of delta-covering problems. In particular, we obtain an efficient algorithm that decides whether an evenly directed 4-regular graph admits a pair of compatible euler tours.

Our results hold for skew-symmetric matrices over any field. However, for fields of characteristic two, we insist that skew-symmetric matrices have diagonal entries zero. We also provide a characterization of binary delta-matroids (these are delta-matroids represented by symmetric matrices over GF(2)). It is shown that a binary delta-matroid is an elementary projection of an even binary delta-matroid (that is, a delta-matroid that can be represented by a skew-symmetric matrix over GF(2)). This characterization enables us to efficiently solve the delta-parity problem for binary delta-matroids.

2. Preliminaries

2.1. Skew-symmetric matrices

A matrix A whose row/column set is indexed by a finite set V is said to be *skew-symmetric* if A is equal to the transpose of $-A$, and all diagonal entries of A are zero. (For fields of characteristic different from two, the condition that A has a zero diagonal is implied by the condition that $A = -A^t$.) The *support graph* of A is the graph $G = (V, E)$ with vertex set V and edge set $E = \{(u, v) \mid A_{uv} \neq 0\}$.

Skew-symmetric matrices have the peculiar property that their determinants are perfect squares. The square root of the determinant is called the *Pfaffian* of A. The

Pfaffian of A can be computed by taking weighted sums over all perfect matchings M of the support graph G:

$$\mathrm{Pf}\, A = \sum_{M} \sigma_M \prod_{(u,v)\in M} A_{uv},$$

where σ_M takes ± 1 in a suitable manner, see [15]. In particular, A is singular if G has no perfect matching (as is the case when $|V|$ is odd). On the other hand, if G has exactly one perfect matching, then A is nonsingular.

A matrix A' is said to be *congruent* to A if there exists a nonsingular matrix Q such that $A' = Q^t AQ$. The operation converting A to A' is called a *congruence transformation*. Note that skew-symmetry and nonsingularity are invariant under congruence transformations. Our augmenting path algorithm, in Section 5, implicitly uses the following fact: *A skew-symmetric matrix A is nonsingular if and only if there exists a matrix A', congruent to A, whose support graph has a unique perfect matching.*

We use the following lemma to show that a graph has a unique perfect matching. (If X is a subset of vertices of a graph G, we let $G[X]$ denote the subgraph of G that is induced by X.)

Lemma 2.1. *Let $G = (W, E)$ be a graph with vertex set W of even cardinality. Suppose the vertex set W can be partitioned into odd sets $U_0, U_1, \ldots, U_h$ and vertices $y_0, y_1, \ldots, y_h$ such that*

(1) *$G[U_i \cup \{y_i\}]$ has a unique perfect matching for each $i = 0, 1, \ldots, h$;*
(2) *$i < j$ implies $(w, z) \notin E$ for $w \in U_i$ and $z \in U_j \cup \{y_j\}$.*

Then $G = (W, E)$ has a unique perfect matching.

Throughout this paper, $A[X]$ denotes a principal submatrix of A indexed by $X \subseteq V$, and $A[X, Y]$ designates a submatrix whose row and column are indexed by X and Y, respectively. We also denote by $X \triangle Y$ the symmetric difference of X and Y, that is, $X \triangle Y = (X \cup Y) - (X \cap Y)$.

Suppose the skew-symmetric matrix A has the following form:

$$\begin{array}{cc} & \begin{array}{cc} Y & X \end{array} \\ A = & \begin{pmatrix} \alpha & \beta \\ -\beta^t & \gamma \end{pmatrix}. \end{array}$$

We define a matrix $A * Y$ by

$$\begin{array}{cc} & \begin{array}{cc} Y & \qquad X \end{array} \\ A * Y = & \begin{pmatrix} \alpha^{-1} & \alpha^{-1}\beta \\ \beta^t\alpha^{-1} & \gamma + \beta^t\alpha^{-1}\beta \end{pmatrix}. \end{array}$$

This operation converting A to $A * Y$ is called a *pivoting*. The following theorem is fundamental to linear delta-matroids.

Theorem 2.1 (Tucker [19]). *Let $A[Y]$ be a nonsingular principal submatrix of a skew-symmetric matrix A. Then, for all $S \subseteq V$,*

$$\det(A * Y)[S] = \det A[Y \triangle S]/\det A[Y].$$

2.2. Delta-matroids

A *delta-matroid* is a pair $(V, \mathscr{F})$ of a finite set V and a nonempty family $\mathscr{F}$ of its subsets, called *feasible sets*, that satisfy the *symmetric exchange axiom*.

(DM) For $F, F' \in \mathscr{F}$ and $x \in F \triangle F'$, there exists $y \in F \triangle F'$ such that $F \triangle \{x, y\} \in \mathscr{F}$.

We call a delta-matroid whose feasible sets all have the same cardinality modulo two an *even delta-matroid*. The base family of a matroid is an even delta-matroid. In fact, a delta-matroid is a matroid if and only if the feasible sets all have the same cardinality. Given a feasible set $F \in \mathscr{F}$ of an even delta-matroid $\mathbf{M} = (V, \mathscr{F})$, we consider a graph $G_F = (V, E_F)$ with an edge set $E_F = \{(u, v) \mid F \triangle \{u, v\} \in \mathscr{F}\}$. The graph G_F is called the *fundamental graph* of $\mathbf{M}$ with respect to F.

For a delta-matroid $\mathbf{M} = (V, \mathscr{F})$ and $X \subseteq V$, we denote $\mathbf{M} \triangle X = (V, \mathscr{F} \triangle X)$, where $\mathscr{F} \triangle X = \{F \triangle X \mid F \in \mathscr{F}\}$. It is easy to see that $\mathbf{M} \triangle X$ is a delta-matroid. This operation is referred to as *twisting* by X, and $\mathbf{M} \triangle X$ is said to be *equivalent* to $\mathbf{M}$. The delta-matroid $\mathbf{M}^* = \mathbf{M} \triangle V$ is the *dual* of $\mathbf{M}$. It is also easy to see that $\mathbf{M} \backslash X = (V \backslash X, \mathscr{F} \backslash X)$ defined by $\mathscr{F} \backslash X = \{F \mid F \in \mathscr{F},\ F \subseteq V \backslash X\}$ is a delta-matroid, provided that $\mathscr{F} \backslash X$ is nonempty. This operation is referred to as the *deletion* of X. The *contraction* of $\mathbf{M}$ by X means $(\mathbf{M} \triangle X) \backslash X$, denoted $\mathbf{M}/X$. Note that evenness is invariant under these operations.

Let A be a skew-symmetric matrix whose row/column set is identified with V. Then the family of the nonsingular principal submatrices $\mathscr{F}(A) = \{X \mid \operatorname{rank} A[X] = |X|\}$ satisfies (DM), and hence $\mathbf{M}(A) = (V, \mathscr{F}(A))$ forms a delta-matroid [2]. We call a delta-matroid *linear* if it is equivalent to $\mathbf{M}(A)$ for some skew-symmetric matrix A. Since every skew-symmetric matrix of odd size is singular, a linear delta-matroid is an even delta-matroid.

For a feasible set F of a linear delta-matroid $\mathbf{M} = \mathbf{M}(A) \triangle X$, we call a skew-symmetric matrix A_F a fundamental matrix of $\mathbf{M}$ with respect to F if it satisfies $\mathbf{M}(A_F) = \mathbf{M} \triangle F$. Note that the support graph of A_F coincides with the fundamental graph G_F. A fundamental matrix can be constructed explicitly. In fact it follows from Theorem 2.1 that $A * Y$ with $Y = X \triangle F$ is a fundamental matrix of $\mathbf{M}$ with respect to F.

For a pair of delta-matroids $\mathbf{M}_1 = (V_1, \mathscr{F}_1)$ and $\mathbf{M}_2 = (V_2, \mathscr{F}_2)$ with $V_1 \cap V_2 = \emptyset$, the *direct sum* $\mathbf{M}_1 \oplus \mathbf{M}_2$ designates the delta-matroid $(V, \mathscr{F})$ defined by $V = V_1 \cup V_2$ and $\mathscr{F} = \{F_1 \cup F_2 \mid F_1 \in \mathscr{F}_1, F_2 \in \mathscr{F}_2\}$. If $\mathbf{M}_1 = \mathbf{M}(A_1) \triangle X_1$ and $\mathbf{M}_2 = \mathbf{M}(A_2) \triangle X_2$, the direct sum $\mathbf{M}_1 \oplus \mathbf{M}_2$ is represented as $\mathbf{M}(A) \triangle X$, where

$X = X_1 \cup X_2$ and

$$A = \begin{array}{c} \\ \\ \end{array}\begin{array}{cc} V_1 & V_2 \\ \end{array}$$

$$A = \begin{pmatrix} A_1 & O \\ O & A_2 \end{pmatrix}.$$

3. The delta-parity problem

Let $\mathbf{M} = (V, \mathscr{F})$ be a delta-matroid on V with $|V|$ even. Consider a partition Π of V into pairs, called *lines*. Each element $v \in V$ has its *mate* $\bar{v}$ such that $\{v, \bar{v}\}$ is a line. For $F \subseteq V$, we denote by $\delta_\Pi(F)$ the number of elements $v \in F$ whose mate $\bar{v}$ is out of F. In other words, $\delta_\Pi(F)$ is the number of lines exactly one element of which belongs to F. Then the delta-parity problem is to find a feasible set $F \in \mathscr{F}$ that minimizes $\delta_\Pi(F)$. We denote by $\delta(\mathbf{M}, \Pi)$ the optimal value of this problem, that is,

$$\delta(\mathbf{M}, \Pi) = \min\{\delta_\Pi(F) \mid F \in \mathscr{F}\}.$$

In the following, we introduce a lower bound on $\delta(\mathbf{M}, \Pi)$. It will be shown later, as the main result of this paper, that the lower bound is tight for linear delta-matroids. This is a natural extension of the Lovász minimax theorem for the linear matroid parity problem.

In order to present the lower bound, we first introduce a strong map operation for delta-matroids. Let $\mathbf{M} = (V, \mathscr{F})$ and $\mathbf{M}^\circ = (V, \mathscr{F}^\circ)$ be delta-matroids. We say that $\mathbf{M}^\circ$ is a *strong map* of $\mathbf{M}$ if there exists Z and a delta-matroid $\mathbf{M}^+ = (V \cup Z, \mathscr{F}^+)$ with $\mathbf{M} = \mathbf{M}^+ \backslash Z$ and $\mathbf{M}^\circ = \mathbf{M}^+/Z$; this is denoted $\mathbf{M} \leftrightarrow \mathbf{M}^\circ$. Note that $\mathbf{M}$ is a strong map of $\mathbf{M}^\circ$ if and only if $\mathbf{M}^\circ$ is a strong map of $\mathbf{M}$. (We do not know of a pair of delta-matroids, on a common ground set, that are not related by a strong map.) However, it is possible that this operation may not be transitive. If $\mathbf{M} \leftrightarrow \mathbf{M}^\circ$, then the *distance* between $\mathbf{M}$ and $\mathbf{M}^\circ$, denoted by $\mathrm{dist}(\mathbf{M}, \mathbf{M}^\circ)$, is the minimum cardinality of Z such that there exists a delta-matroid $\mathbf{M}^+ = (V \cup Z, \mathscr{F}^+)$ with $\mathbf{M} = \mathbf{M}^+ \backslash Z$ and $\mathbf{M}^\circ = \mathbf{M}^+/Z$.

The following result provides a lower bound on the optimal value of the delta-parity problem.

Lemma 3.1. *For a pair of delta-matroids* $\mathbf{M}$ *and* $\mathbf{M}^\circ$ *related by a strong map, we have* $\delta(\mathbf{M}, \Pi) \geqslant \delta(\mathbf{M}^\circ, \Pi) - \mathrm{dist}(\mathbf{M}, \mathbf{M}^\circ)$.

Proof. If $\mathrm{dist}(\mathbf{M}, \mathbf{M}^\circ) \leqslant \ell$, there exists a sequence $\mathbf{M}_0, \mathbf{M}_1, \ldots, \mathbf{M}_\ell$ of delta-matroids with $\mathbf{M}_0 = \mathbf{M}$, $\mathbf{M}_\ell = \mathbf{M}^\circ$, and $\mathrm{dist}(\mathbf{M}_{i-1}, \mathbf{M}_i) = 1$ for $i = 1, \ldots, \ell$. Hence, it suffices to show the statement in the case of $\mathrm{dist}(\mathbf{M}, \mathbf{M}^\circ) = 1$. Let $\mathbf{M}^+ = (V \cup \{z\}, \mathscr{F}^+)$ be a delta-matroid with $\mathbf{M} = \mathbf{M}^+ \backslash \{z\}$ and $\mathbf{M}^\circ = \mathbf{M}^+/\{z\}$. Then there exists a feasible set $F^+ \in \mathscr{F}^+$ that contains z. Suppose $F \in \mathscr{F}$ is an optimal solution of the delta-parity problem for $\mathbf{M}$. The symmetric exchange axiom (DM) implies that there exists $y \in F \triangle F^+$ such that $F^\bullet = F \triangle \{y, z\} \in \mathscr{F}^+$. Since $z \in F^\bullet$, we have $F^\circ = F^\bullet - \{z\} \in \mathscr{F}^\circ$.

Then $\delta(\mathbf{M},\Pi)=\delta_\Pi(F)\geqslant\delta_\Pi(F^\circ)-1\geqslant\delta(\mathbf{M}^\circ,\Pi)-1$, which completes the proof. □

We now exploit properties of even delta-matroids. Suppose $\mathbf{M}=(V,\mathscr{F})$ is an even delta-matroid, and consider a direct sum decomposition $\mathbf{M}=\mathbf{M}_1\oplus\cdots\oplus\mathbf{M}_\ell$. Note that each component $\mathbf{M}_i$ is also an even delta-matroid. If every feasible set of $\mathbf{M}_i$ is of odd cardinality, $\mathbf{M}_i$ is called an odd component. We denote by $\mathrm{odd}(\mathbf{M},\Pi)$ the number of odd components in the finest direct sum decomposition which is compatible with the partition Π (that is, the ground set of each $\mathbf{M}_i$ is a union of lines).

Lemma 3.2. *An even delta-matroid* $\mathbf{M}$ *satisfies* $\delta(\mathbf{M},\Pi)\geqslant\mathrm{odd}(\mathbf{M},\Pi)$.

Combining this with Lemma 3.1, we obtain the following lemma.

Lemma 3.3. *Suppose* $\mathbf{M}$ *and* $\mathbf{M}^\circ$ *are even delta-matroids related by a strong map. Then* $\delta(\mathbf{M},\Pi)\geqslant\mathrm{odd}(\mathbf{M}^\circ,\Pi)-\mathrm{dist}(\mathbf{M},\mathbf{M}^\circ)$.

This lower bound on $\delta(\mathbf{M},\Pi)$ is in fact tight for a suitable $\mathbf{M}^\circ$ in the case of linear delta-matroids. The precise statement is as follows.

Theorem 3.1. *For a linear delta-matroid* $\mathbf{M}$, *we have*

$$\delta(\mathbf{M},\Pi)=\max\{\mathrm{odd}(\mathbf{M}^\circ,\Pi)-\mathrm{dist}(\mathbf{M},\mathbf{M}^\circ)\mid\mathbf{M}\leftrightarrow\mathbf{M}^\circ,\mathbf{M}^\circ\text{: } \textit{linear}\}. \tag{1}$$

We will prove this theorem via an augmenting path algorithm for the delta-parity problem on linear delta-matroids.

4. Applications and related problems

In applications, Theorem 3.1 does not always provide a satisfactory minimax theorem. Typically, $\mathbf{M}$ is in some restricted class of delta-matroids and we would like $\mathbf{M}^\circ$ to also be chosen from the same class. Often, by closer analysis of the algorithm, we can place such conditions on $\mathbf{M}^\circ$. For instance, if $\mathbf{M}$ is equivalent to a matroid, then so is $\mathbf{M}^\circ$. Also, if $\mathbf{M}=\mathbf{M}_1\oplus\mathbf{M}_2$ then $\mathbf{M}^\circ={\mathbf{M}^\circ}_1\oplus{\mathbf{M}^\circ}_2$, where $\mathbf{M}_1\leftrightarrow{\mathbf{M}^\circ}_1$ and $\mathbf{M}_2\leftrightarrow{\mathbf{M}^\circ}_2$.

4.1. Linear matroid parity

Given a matroid $\mathbf{M}=(V,\mathscr{B})$ with the base family $\mathscr{B}$ and a partition Π of V into pairs, the *matroid parity problem* is to find a base containing the maximum number of lines (or equivalently to find a largest independent set containing only lines). Let $v(\mathbf{M},\Pi)$ denote the optimal value of the matroid parity problem. Then it is obvious that $2v(\mathbf{M},\Pi)=\mathrm{rank}\,\mathbf{M}-\delta(\mathbf{M},\Pi)$.

Thus, it is clear that the delta-matroid parity problem is a generalization of the matroid parity problem. However, it is not immediately clear that linear matroids are linear delta-matroids; this is outlined below.

Let N be a matrix whose columns are indexed by V, and let $\mathscr{B}$ be the set of all sets that index maximal linearly independent sets of columns of N. Then, $\mathbf{M} = (V, \mathscr{B})$ forms a matroid which is called *linear*. Let B be a basis of $\mathbf{M}$. By elementary row operations, we may assume that N has the following form:

$$N = \begin{array}{cc} B & V\backslash B \\ (I & N') \end{array},$$

where I is a B by B identity matrix. For $F \subseteq V$, $F \in \mathscr{B}$ if and only if $N'[B\backslash F, F\backslash B]$ is square and nonsingular. We now define a V by V skew-symmetric matrix A as follows:

$$A = \begin{array}{c} \begin{array}{cc} B & V\backslash B \end{array} \\ \begin{pmatrix} O & N' \\ -(N')^t & O \end{pmatrix} \end{array}.$$

For $S \subseteq V$, $A[S \triangle B]$ is nonsingular if and only if $N'[B\backslash S, S\backslash B]$ is square and nonsingular. Therefore, we have

$$\mathbf{M} = \mathbf{M}(A) \triangle B.$$

Thus, linear matroids are indeed linear delta-matroids.

Lovász [13,14] proved a minimax theorem similar to Theorem 3.1. While this does not follow immediately from Theorem 3.1, it can be proved using additional structure from the algorithm. Indeed, this is Gabow and Stallmann's proof [9] of Lovász's minimax theorem.

4.2. The delta-covering problem

Given a pair of delta-matroids $(V, \mathscr{F}_1)$ and $(V, \mathscr{F}_2)$, find $F_1 \in \mathscr{F}_1$ and $F_2 \in \mathscr{F}_2$ maximizing $|F_1 \triangle F_2|$. This is the *delta-covering problem* posed by Bouchet [3]. The delta-covering problem is closely related to two natural decision problems.

Partition problem: Given a pair of delta-matroids $(V, \mathscr{F}_1)$ and $(V, \mathscr{F}_2)$, does there exist a partition (F_1, F_2) of V such that $F_1 \in \mathscr{F}_1$ and $F_2 \in \mathscr{F}_2$?

Intersection problem: Given a pair of delta-matroids $\mathbf{M}_1$ and $\mathbf{M}_2$, do $\mathbf{M}_1$ and $\mathbf{M}_2$ share a common feasible set?

The connection between the delta-covering problem and the partition problem is clear. The intersection problem for $\mathbf{M}_1$ and $\mathbf{M}_2$ is the partition problem for $\mathbf{M}_1$ and $\mathbf{M}_2 \triangle V$. Given its relation to these very familiar problems, the delta-covering problem may seem more natural than the delta-parity problem. However, we shall see that the delta-covering problem and the delta-parity problem are equivalent.

The delta-parity problem is a special case of delta-covering problems as follows. Let $\mathscr{L}$ denote the family of subsets of V that can be represented as a union of lines. Then $(V, \mathscr{L})$ forms an even delta-matroid, and $\delta_\Pi(F) = |V| - \max\{|F \triangle L| \mid L \in \mathscr{L}\}$

holds for $F \subseteq V$. Hence, the delta-parity problem is a delta-covering problem on $(V, \mathscr{F})$ and $(V, \mathscr{L})$.

On the other hand, an arbitrary delta-covering problem can be reduced to a delta-parity problem. Given a pair of delta-matroids $(V, \mathscr{F}_1)$ and $(V, \mathscr{F}_2)$, denote their copies by $\mathbf{M}_1 = (V_1, \mathscr{F}'_1)$ and $\mathbf{M}_2 = (V_2, \mathscr{F}'_2)$, respectively. Let $\mathbf{M} = (V_1 \cup V_2, \mathscr{F})$ be the direct sum of $\mathbf{M}_1$ and $\mathbf{M}_2^*$, the dual of $\mathbf{M}_2$, and Π be the partition of the ground set $V_1 \cup V_2$ into the pairs each of which corresponds to an original element in V. For a pair of feasible sets $F_1 \in \mathscr{F}_1$ and $F_2 \in \mathscr{F}_2$, it is easy to see that $\delta_\Pi(F) = |V| - |F_1 \triangle F_2|$ holds for $F = F'_1 \cup (V_2 - F'_2) \in \mathscr{F}$, where F'_1 and F'_2 are the copies of F_1 and F_2. Therefore, the delta-covering problem on $(V, \mathscr{F}_1)$ and $(V, \mathscr{F}_2)$ coincides with the delta-parity problem for $\mathbf{M}$ with the partition Π.

It is not difficult to convert Theorem 3.1 to a minimax relation on the delta-covering problem for linear delta-matroids. For a pair of even delta-matroids $\mathbf{M}_1 = (V, \mathscr{F}_1)$ and $\mathbf{M}_2 = (V, \mathscr{F}_2)$, consider the finest partition $(V_1, \dots, V_\ell)$ of V that simultaneously gives direct-sum decompositions of both $\mathbf{M}_1$ and $\mathbf{M}_2$. We say V_i is an odd component with respect to $(\mathbf{M}_1, \mathbf{M}_2)$ if $|F_1 \cap V_i| + |F_2 \cap V_i| - |V_i|$ is odd for a pair of feasible sets $F_1 \in \mathscr{F}_1$ and $F_2 \in \mathscr{F}_2$. We denote by $\mathrm{odd}(\mathbf{M}_1, \mathbf{M}_2)$ the number of odd components with respect to $(\mathbf{M}_1, \mathbf{M}_2)$. Then we have the following corollary. (Again, this is not a direct corollary of Theorem 3.1, we use some additional properties from the algorithm.)

Corollary 4.1. *For a pair of linear delta-matroids* $\mathbf{M}_1 = (V, \mathscr{F}_1)$ *and* $\mathbf{M}_2 = (V, \mathscr{F}_2)$, *we have*

$$\begin{aligned}&\max\{|F_1 \triangle F_2| \mid F_1 \in \mathscr{F}_1, F_2 \in \mathscr{F}_2\} \\ &\quad = |V| - \max\{\mathrm{odd}(\mathbf{M}_1{}^\circ, \mathbf{M}_2{}^\circ) - \mathrm{dist}(\mathbf{M}_1, \mathbf{M}_1{}^\circ) - \mathrm{dist}(\mathbf{M}_2, \mathbf{M}_2{}^\circ)\},\end{aligned}$$

where $\mathbf{M}_1{}^\circ \leftrightarrow \mathbf{M}_1$ *and* $\mathbf{M}_2{}^\circ \leftrightarrow \mathbf{M}_2$, *and* $\mathbf{M}_1{}^\circ$ *and* $\mathbf{M}_2{}^\circ$ *are chosen among linear delta-matroids.*

4.3. The partition problem and compatible Euler tours

The following problem is an important special case of the partition problem (stated above): *Given a delta-matroid* $\mathbf{M} = (V, \mathscr{F})$, *can* V *be partitioned into two feasible sets of* $\mathbf{M}$? There is some hope that there may exist an efficient algorithm for solving this problem for any even delta-matroid. However, even for linear delta-matroids the only known algorithm is that presented here. This partition problem, for a linear delta-matroid $\mathbf{M}(A) \triangle X$, is the problem of deciding if there exists a partition (F_1, F_2) of V such that $A[F_1]$ and $A[F_2]$ are both nonsingular. (The analogous problem for a symmetric matrix A remains open.) This partition problem also contains an interesting graph theoretic problem.

Let $\vec{G} = (V, \vec{E})$ be a connected directed graph such that each vertex has two entering and two leaving arcs. We are interested in the euler tours in $\vec{G}$. An euler tour induces a pairing of the entering arcs of a vertex with its leaving arcs. There are two possible such pairings at each vertex; we refer to any such pairing as a *bitransition*. Two euler tours are *compatible* if they do not use a common bitransition. We are interested in the problem of deciding whether there exists a pair of compatible euler tours. This problem was proposed by Bouchet [3], and provided much of the motivation for studying the delta-covering problem.

We fix a particular *reference* bitransition at each vertex, the collection of which we denote by R. Note that R need not correspond to an euler tour, as it may contain a number of disconnected walks. We encode an euler tour T, of $\vec{G}$, by the set $R(T)$ of vertices at which R and T share common bitransitions. Clearly, two euler tours T_1 and T_2 are compatible if and only if $(R(T_1), R(T_2))$ partitions V. Now let $\mathscr{F}$ be the collection of sets $R(T)$ taken over all euler tours T of $\vec{G}$. We call $(V, \mathscr{F})$ an *eulerian delta-matroid*; eulerian delta-matroids are in fact linear delta-matroids over an arbitrary field, see Bouchet [3]. The partition problem for this delta-matroid is the problem of finding compatible euler tours in $\vec{G}$.

Our augmenting path algorithm for the delta-parity problem provides an efficient algorithm for finding compatible euler tours. Unfortunately, we have been unable to deduce a nice graphic condition for the existence of a pair of compatible euler tours from Theorem 3.1. It would be interesting if, for an eulerian delta-matroid, Theorem 3.1 still holds when $\mathbf{M}^{\circ}$ is chosen among eulerian delta-matroids.

4.4. Binary delta-matroids

Recall that we have defined skew-symmetric matrices so that, even over fields of characteristic two, we have a zero diagonal. This definition ensures that linear delta-matroids are even. However, there are alternative notions of "linear delta-matroids." For a symmetric matrix A, $\mathbf{M}(A)$ is a delta-matroid (see [2,5]). The delta-parity problem is also of some interest for delta-matroids arising from symmetric matrices. We shall describe a trick that enables us to solve the delta-parity problem for symmetric binary matrices.

A *binary delta-matroid* is a delta-matroid that is equivalent to $\mathbf{M}(A) = (V, \mathscr{F}(A))$ for some symmetric matrix A over GF(2). Recall $\mathscr{F}(A) = \{X \mid \operatorname{rank} A[X] = |X|\}$, where the row/column set of A is identified with V. A binary delta-matroid is not necessarily even. In fact, a binary delta-matroid is even if and only if it is representable by a skew-symmetric matrix over GF(2).

For a delta-matroid $\mathbf{M} = (V, \mathscr{F})$ and $X \subseteq V$, we define $\mathscr{F}|X = \{F \backslash X \mid F \in \mathscr{F}\}$. Then it can be shown that $\mathbf{M}|X = (V \backslash X, \mathscr{F}|X)$ is a delta-matroid. This operation, or $\mathbf{M}|X$ itself, is called the *projection* of $\mathbf{M}$ on X. In particular, we call $\mathbf{M}|X$ an *elementary projection* if X is a singleton. Note that projection does not necessarily preserve evenness.

Lemma 4.1. *Every binary delta-matroid is an elementary projection of some even binary delta-matroid.*

Proof. Suppose A is a symmetric matrix over GF(2) with row/column set V. Let ψ denote a vector corresponding to the diagonal of A, i.e., $\psi_v = A_{vv}$ for $v \in V$. Consider two matrices

$$\hat{A} = \begin{pmatrix} A+\Psi & \psi \\ \psi^t & 1 \end{pmatrix}, \quad \tilde{A} = \begin{pmatrix} A+\Psi & \psi \\ \psi^t & 0 \end{pmatrix},$$

where $\Psi = \psi\psi^t$ and the row/column sets of $\hat{A}$ and $\tilde{A}$ are both $V \cup \{u\}$. For $Y \subseteq V$, we have $\det A[Y] = \det \hat{A}[Y \cup \{u\}] = \det \tilde{A}[Y] + \det \tilde{A}[Y \cup \{u\}]$. Since $\tilde{A}$ is skew-symmetric, $\tilde{A}[Y]$ or $\tilde{A}[Y \cup \{u\}]$ is singular. Hence, $A[Y]$ is nonsingular if and only if either $\tilde{A}[Y]$ or $\tilde{A}[Y \cup \{u\}]$ is nonsingular. Thus we have $\mathbf{M}(A) = \mathbf{M}(\tilde{A})|\{u\}$, which implies $\mathbf{M}(A) \triangle X = (\mathbf{M}(\tilde{A}) \triangle X) \mid \{u\}$ for $X \subseteq V$. $\square$

It is obvious from the proof that the converse also holds. Namely, an elementary projection of an even binary delta-matroid is a binary delta-matroid. It should be remarked, however, that projection does not necessarily preserve binary representability.

We now discuss how to solve the delta-parity problem for a noneven binary delta-matroid $\mathbf{M} = (V, \mathscr{F})$ represented as $\mathbf{M}(A) \triangle X$ for $X \subseteq V$. Consider the even binary delta-matroid $\mathbf{M}' = \mathbf{M}(\tilde{A}) \triangle X$ constructed as in the proof of Lemma 4.1. Since $\mathbf{M} = \mathbf{M}'|\{u\}$, a feasible set of $\mathbf{M}$ is either a feasible set of $\mathbf{M}' \backslash \{u\}$ or $\mathbf{M}'/\{u\}$. Note that both of these delta-matroids are even and binary. Therefore, in order to solve the delta-parity problem for $\mathbf{M}$, we need only apply the augmenting path algorithm to $\mathbf{M}' \backslash \{u\}$ and $\mathbf{M}'/\{u\}$ independently and compare the solutions. The better of the two gives an optimal solution for $\mathbf{M}$.

It should be remarked that the above argument, including Lemma 4.1, is valid not only over GF(2) but also over an arbitrary finite field of characteristic two. The vector ψ in the proof of Lemma 4.1 should be defined in such a way that ψ_v is a square root of A_{vv}. Hence, we can efficiently solve the delta-parity problem for delta-matroids represented by symmetric matrices over a finite field of characteristic two.

5. An augmenting path algorithm

In this section we present an augmenting path algorithm for solving the delta-parity problem on linearly represented delta-matroids. The algorithm is relatively straightforward, but the proof of correctness is, as should be expected, rather technical.

Let $\mathbf{M} = (V, \mathscr{F})$ be an even delta-matroid represented as $\mathbf{M} = \mathbf{M}(A) \triangle X$ by a skew-symmetric matrix A and $X \subseteq V$, let Π be a partition of V into pairs, and let $F \in \mathscr{F}$. We will describe an algorithm that either concludes, with the aid of Lemma 3.3, that F is an optimal solution for the delta-parity problem, or finds a feasible set

F' such that $\delta_\Pi(F') = \delta_\Pi(F) - 2$. One may then solve the delta-parity problem by beginning with the feasible set X and repeatedly applying the algorithm.

The search procedure for an augmenting path with respect to $F \in \mathscr{F}$ works primarily on a skew-symmetric matrix $A^\sharp$ whose row/column set is denoted by $V^\sharp$. The algorithm initially sets $A^\sharp$ to the fundamental matrix A_F and augments $A^\sharp$ each time it generates a new element to be added to $V^\sharp$. The new elements are called *transforms*. We denote by $\tau(p,q,b)$ a transform generated with reference to three elements $p, q \in V$ and $b \in V^\sharp$. Whenever the algorithm generates a transform $\tau(p,q,b)$, both $A^\sharp_{pb}$ and $A^\sharp_{qb}$ are nonzero. On generating a transform $t = \tau(p,q,b)$, the algorithm computes $\mu = A^\sharp_{bp}A^\sharp_{*q} - A^\sharp_{bq}A^\sharp_{*p}$ (where $A^\sharp_{*v}$ denotes the column of $A^\sharp$ indexed by v), adds t to $V^\sharp$, and updates $A^\sharp$ as

$$A^\sharp \leftarrow \begin{array}{c} \begin{array}{cc} & t \end{array} \\ \begin{pmatrix} A^\sharp & \mu \\ -\mu^t & 0 \end{pmatrix} \end{array}.$$

The vector μ is chosen to be a linear combination of the columns $A^\sharp_{*q}$ and $A^\sharp_{*p}$ such that $\mu_b = 0$.

The algorithm uses the support graph $G^\sharp = (V^\sharp, E^\sharp)$ of $A^\sharp$. Recall that $\delta_\Pi(F)$ is the number of lines that have exactly one element in F. Such lines are called *source lines* and we denote by S the set of elements that belong to source lines.

For each vertex $v \in V$, let $v^\natural$ denote p if v is a transform $\tau(p,q,b)$, and v otherwise. A *parity path* is a sequence of vertices $v_0v_1\cdots v_l$ such that $v_i^\natural$ are all distinct, $(v_{i-1}, v_i) \in E^\sharp$ for odd i and $\{v_{i-1}^\natural, v_i^\natural\}$ is a line of Π for even i. (Note that a parity path is a sequence of vertices that does not necessarily correspond to a path in $G^\sharp$.) For a parity path P, we denote by $\bar{P}$ the reverse sequence of P. Also, we let $P^\natural = \{v^\natural \mid v \in P\}$. A parity path $P = v_0v_1\cdots v_l$ is said to be an *augmenting path* if it satisfies the following conditions:

(AP1) l is odd;
(AP2) $v_0, v_l \in S$, but $\{v_0, v_l\}$ is not a line;
(AP3) $v_i^\natural \notin S$ for $i = 1, \ldots, l-1$;
(AP4) $\tau(p,q,b) \in P$ implies $q \in P$ and $p \notin P$;
(AP5) $M = \{(v_{i-1}, v_i) \mid i\text{: odd}\}$ is the only one perfect matching in $G^\sharp[P]$.

After finding an augmenting path P, the algorithm updates the feasible solution F to $F \triangle P^\natural$. It is obvious from (AP1)–(AP4) that $\delta_\Pi(F \triangle P^\natural) = \delta_\Pi(F) - 2$. The following lemma guarantees that $F \triangle P^\natural \in \mathscr{F}$.

Lemma 5.1. *If P is an augmenting path, $A_F[P^\natural]$ is nonsingular.*

Proof. Since $A^\sharp[P]$ is congruent to $A_F[P^\natural]$ by (AP4), it suffices to show that $A^\sharp[P]$ is nonsingular. A perfect matching of $G^\sharp[P]$ corresponds to a nonzero term of the

Pfaffian of $A^\sharp[P]$, the square of which is equal to the determinant of $A^\sharp[P]$. Hence the uniqueness of the perfect matching implies that $A^\sharp[P]$ is nonsingular. □

We now describe the search procedure for finding an augmenting path with respect to a feasible set. Basically, it performs a breadth-first search using a queue to grow parity paths of odd cardinality from vertices in S. A vertex v is *labeled* and put into the queue when it is reached by a parity path (or *search path*) $P(v)$. Once labeled, a vertex is never unlabeled nor relabeled. The algorithm picks the first labeled element, say v, from the queue, sequentially assigns a number $\sigma(v)$ to v, and then examines its neighbors. This process is said to *scan* v.

The algorithm associates a set $B(x) \subseteq V^\sharp$ with every labeled vertex $x \in V^\sharp$; once defined $B(x)$ is never modified. We call $B(x)$ a *blossom*. A blossom is a union of lines and possibly transforms. Source lines in particular belong to distinct blossoms. A blossom free from source lines is accompanied by a vertex called the *bud*, while a blossom including a source line has no bud. A vertex in a blossom is called a *tip* of the blossom if it is adjacent to the bud. Blossoms may be nested or disjoint but never cross; that is, $B(x) \subseteq B(y)$, $B(y) \subseteq B(x)$, or $B(x) \cap B(y) = \emptyset$ for any labeled x and y. For any $x \in V^\sharp$, we denote by $B^\bullet(x)$ the maximal blossom that contains x. Note that maximal blossoms are disjoint. It is helpful to imagine that each maximal blossom has been shrunk to a line that consists of two supernodes: one for the tips and the other for the rest.

It is possible to implement the algorithm without keeping $P(v)$ and $B(x)$ explicitly, which is important from the viewpoint of computational efficiency. Gabow and Stallmann [9] discuss this issue in greater detail for the linear matroid parity problem.

Recall that the algorithm works with a graph $G^\sharp = (V^\sharp, E^\sharp)$, which is implicit in the description below. For $x \in P(v)$, we let $P(v|x)$ denote the subsequence of $P(v)$ from the successor of x to v.

Algorithm. Search(F)

Step 0: Let S be the set of elements in the source lines. For every $s \in S$, set $B(s) := \{s, \bar{s}\}$, label s with $P(s) := s$, and put s into the queue. Put $V^\sharp := V$, $A^\sharp := A_F$, and $\kappa := 1$.

Step 1: If the queue is empty, then stop. Otherwise, pick the first element v from the queue, assign $\sigma(v) := \kappa$.

Step 2: For each scanned vertex u adjacent to v, if $u \notin B^\bullet(v)$, apply Blossom(v, u) choosing u in the increasing order of $\sigma(u)$.

Step 3: For each unlabeled vertex u adjacent to v, if u does not belong to any blossoms, put $\lambda(u) := \sigma(v)$ and do the following steps.

(3-1): If the mate $\bar{u}$ is also adjacent to v, then create a transform $t = \tau(\bar{u}, u, v)$, label it with $P(t) := P(v)ut$, set $B(t) := \{u, \bar{u}, t\}$, and put t into the queue. Augment $V^\sharp$ and $A^\sharp$ accordingly.

(3-2): If the mate $\bar{u}$ is not adjacent to v, then label $\bar{u}$ with $P(\bar{u}) := P(v)u\bar{u}$, set $B(\bar{u}) := \{u, \bar{u}\}$, and put $\bar{u}$ into the queue.

Step 4: Mark v as scanned, increment κ, and return to Step 1.

Procedure. $\mathsf{Blossom}(v,u)$

Step 0: If the first elements in $P(v)$ and in $P(u)$ belong to different source lines, then terminate with the augmenting path $P := P(v)\overline{P(u)}$.

Step 1: Let b be the last element of $P(v)$ such that $B^\bullet(b) \cap P(u) \neq \emptyset$, and d be the last element of $P(u)$ such that $d \in B^\bullet(b)$. If $d = b$, then put $T := \{p, q\}$ with p and q being the immediate successors of b, respectively, in $P(v)$ and in $P(u)$, create a transform $t = \tau(p,q,b)$, augment $V^\sharp$ and $A^\sharp$ accordingly, and put $B := \{t\}$. Otherwise, put $T := \emptyset$ and $B := B^\bullet(b)$.

Step 2: For each labeled element $x \in P(v|b) \cup P(u|d) - T$, put $B := B \cup B^\bullet(x)$.

Step 3: For each unlabeled element $x \in P(v|b) \cup P(u|d) - T$, label x with

$$P(x) := \begin{cases} P(v)\overline{P(u|x)}x & \text{if } x \in P(u|d), \\ P(u)\overline{P(v|x)}x & \text{if } x \in P(v|b), \end{cases}$$

set $B(x) := B$, and put x into the queue, choosing x in the decreasing order of $\lambda(x)$.

Step 4: If $d = b$, then label t with $P(t) := P(u)\overline{P(v|p)}t$, set $B(t) := B$, and put t into the queue.

The procedure $\mathsf{Blossom}(v,u)$ yields a new blossom. If $b = d$, the vertex b is defined to be the bud of the new blossom. Hence, p and q, possibly among others, are tips of the new blossom. Otherwise (i.e., if $d \neq b$), the new blossom has the same bud as $B^\bullet(b)$, provided that $B^\bullet(b)$ is free from source lines. Step 3 of $\mathsf{Search}(F)$ also yields a new blossom, called a *degenerate blossom*, where v is defined to be the bud. Note that the bud of a blossom is located outside the blossom.

Claim 5.1. *The following conditions hold for any labeled vertex x at* Step 1 *of* $\mathsf{Search}(F)$ *as well as at each termination of the procedure* $\mathsf{Blossom}(v,u)$.

(1) *If $B(x)$ is a blossom with a bud b, then b is scanned and $P(x|b)$ is contained in $B(x)$. Moreover, the first vertex of $P(x|b)$ is a tip of $B(x)$ and the others are labeled.*
(2) *If $B(x)$ is a blossom without a bud, then $P(x)$ is contained in $B(x)$.*

Proof. If $B(x)$ is a degenerate blossom, then $B(x)$ has a bud and (1) is obvious by Step 3 of $\mathsf{Search}(F)$. For the nondegenerate case, suppose inductively that the statements hold before the algorithm performs $\mathsf{Blossom}(v,u)$, now consider applying $\mathsf{Blossom}(v,u)$. If $d \neq b$, then, by the inductive assumption, the statements hold for v and for the vertices labeled in $\mathsf{Blossom}(v,u)$. Hence (2) holds, and to show (1) it suffices to consider the case that $d = b$.

Since b is the last element of $P(v)$ such that $B^\bullet(b) \cap P(u) \neq \emptyset$, the inductive assumption implies that $P(v) = P(b)P(v|b)$ and that b scanned. Since d is the last element of $P(u)$ such that $d \in B^\bullet(b)$, the inductive assumption implies that $P(u) = P(b)P(u|b)$. Therefore, for each vertex x that gets labeled in $\mathsf{Blossom}(v,u)$, we have $P(x|b) \subseteq B(x)$.

The immediate successor of b in $P(v)$ (and in $P(u)$) is a tip of $B(x)$. Thus, the vertex of $P(x|b)$ that succeeds b is a tip of $B(x)$, and, by the inductive assumption, the other vertices $P(x|b)$ are labeled. $\square$

The following result is a straightforward consequence of Claim 5.1 and the description of the algorithm.

Claim 5.2. *If v is a labeled vertex, and $z \in P(v)$, where $z = \tau(p, q, b)$ is a transform, then $q \in P(v)$ and $p \notin P(v)$.*

The following two results are also straightforward.

Claim 5.3. *If v is a labeled vertex and $x \in P(v)$ is not labeled, then $\lambda(x)$ is defined.*

Claim 5.4. *If v is scanned, $y \in P(v)$, and $\lambda(y)$ is defined, then $\sigma(v) \geqslant \lambda(y)$.*

The following result ensures that the search path has odd cardinality.

Claim 5.5. *Let v be a labeled vertex with $P(v) = v_0 v_1, \ldots, v_l$, where $v = v_l$. Then l is even. Moreover, if v_j was not labeled when we labeled v then*

(1) *j is odd, $P(v_{j-1}) = v_0 v_1, \ldots, v_{j-1}$, and $\lambda(v_j) = \sigma(v_{j-1})$ and*
(2) *if $\lambda(v_k)$ is defined and $j \leqslant k \leqslant l$ then $\lambda(v_k) \geqslant \lambda(v_j)$.*

Proof. Suppose inductively that the statements held before we label v. First consider the case that v is labeled in Step 3 of $\mathsf{Search}(F)$. Then v_{l-1} is unlabeled. Moreover, $P(v_{l-2}) = v_0 v_1 \ldots v_{l-2}$, and $\lambda(v_{l-1}) = \sigma(v_{l-2})$. By the inductive assumption, $l-2$ is even, and so is l. Hence (1) and (2) hold when $j = l-1$. Now suppose that $j < l-1$. Since v_j was not labeled when we labeled v, it is not labeled when we labeled v_{l-2}. So we can apply induction to v_j and $P(v_{l-2})$. Thus (1) holds, and (2) holds for $k \leqslant l-2$. Since $\lambda(v_{l-1})$ is defined, and $v^\natural$ is the mate of v_{l-1}, $\lambda(v)$ is not defined. Finally, by Claim 5.4, $\lambda(v_{l-1}) = \sigma(v_{l-2}) \geqslant \lambda(v_j)$.

Now suppose that v is labeled in $\mathsf{Blossom}(r, u)$. Hence, either $v \in P(r)$, $v \in P(u)$, or v is a transform. In each case, the proofs are much the same, so we assume that $v \in P(u)$. Therefore, $P(v) = P(r)\overline{P(u|v)}v$. By the inductive assumption, $P(r)$ is of odd cardinality.

Note that, v was not labeled when we labeled u. Applying induction to v and $P(u)$, we see that $P(u|v)$ is of odd cardinality, which implies l is even. We also see that, for each $z \in P(u|v)$, if $\lambda(z)$ is defined, then $\lambda(z) \geqslant \lambda(v)$. Therefore, each of the vertices in $P(u|v)$ is labeled when we label v. In particular, $v_j \notin P(u|v)$. Therefore, $v_j \in P(r)$. Moreover, v_j was not labeled when we labeled r. Therefore, we can apply induction to v_j and $P(r)$. Thus, (1) holds. Moreover, for each $z \in P(r|v_j)$, if $\lambda(z)$ is defined then $\lambda(z) \geqslant \lambda(v_j)$. Now, to complete the proof of (2), it suffices to prove that $\lambda(v) \geqslant \lambda(v_j)$, since $\lambda(z) \geqslant \lambda(v)$ for each $z \in P(u|v)$ where $\lambda(z)$ is defined.

Note that either $v_j \in B(v)$ or $v_j \in P(u)$. Suppose that $v_j \in P(u)$. However, v_j was not labeled when we labeled u, and $v \in P(u|v_j)$. Therefore, applying induction to v_j and $P(u)$, we get $\lambda(v) \geqslant \lambda(v_j)$ as required. Now suppose that $v_j \in B(v)$. If v_j is labeled in $\mathsf{Blossom}(r, u)$ then, since v_j is labeled after v, we have $\lambda(v) \geqslant \lambda(v_j)$, as required. So we may suppose that v_j is not labeled in Blossom(r, u), and, hence, v_j is a tip of $B(v)$. So, v_{j-1} is the bud of $B(v)$, and, hence $v_{j-1} \in P(u)$. Let w be the successor of v_{j-1} in $P(u)$. So w is also a tip of $B(v)$, and hence was unlabeled when we labeled u. Applying induction to w and $P(u)$, we get $\lambda(v) \geqslant \lambda(w)$ and $\lambda(w) = \sigma(v_{j-1})$. Hence $\lambda(v_j) = \sigma(v_{j-1}) \leqslant \lambda(v)$, as required. □

Suppose that v is a labeled vertex with $P(v) = v_0 v_1 \dots v_l$, where $v = v_l$. We associate a sequence $L(v)$ of labeled vertices with v as follows. If every vertex in $P(v)$ was labeled no later than v, we put $L(v) = P(v)$. Otherwise, we set $L(v) = P(v|v_i)$, where v_i is the last vertex of $P(v)$ that was not labeled when we labeled v. Thus, by Claim 5.5, $L(v)$ has odd cardinality, $P(v) = P(v_{i-1}) v_i L(v)$, and $\sigma(v_{i-1}) = \lambda(v_i)$. Moreover, if v is scanned, then $\sigma(v_{i-1}) < \sigma(v)$. Therefore, the search path $P(v)$ uniquely decomposes as

$$P(v) = L(x_1) y_1 \cdots L(x_k) y_k L(v),$$

where $\sigma(x_1) = \lambda(y_1) < \cdots < \sigma(x_k) = \lambda(y_k)$.

Claim 5.6. *If v is scanned, then, for each $y \in L(v)$, y is scanned and $\sigma(y) \leqslant \sigma(v)$.*

Proof. When a vertex is labeled it is added to the queue, where it is only removed to be scanned. Now v is the last vertex in $L(v)$ to be labeled, so it is also the last to be scanned. □

Claim 5.7. *If v is a labeled vertex with $P(v) = P(x) y L(v)$, $z \in L(v)$ is a transform, and b is the bud of $B(z)$, then $\sigma(b) \geqslant \lambda(y)$.*

Proof. Suppose that $z = \tau(p, q, b)$. Then, $P(z) = P(b) q L(z)$, $\lambda(q) = \sigma(b)$, and, by Claim 5.2, $q \in P(v)$. In fact, either $q = y$ or $q \in L(v)$. If $q \in L(v)$ then, by Claim 5.5, $\lambda(q) \geqslant \lambda(y)$. In either case, we get $\sigma(b) = \lambda(q) \geqslant \lambda(y)$, as required. □

The following result is straightforward.

Claim 5.8. *Let w be a scanned vertex, $v \notin S$ be a vertex that is not a transform, and $(w, v) \in E^\sharp$. Then one of $\lambda(v)$ and $\lambda(\bar{v})$ is defined. Moreover,*

(1) *if $\lambda(v)$ is defined, then $\lambda(v) \leqslant \sigma(w)$ and*
(2) *if $\lambda(\bar{v})$ is defined, then $\lambda(\bar{v}) \leqslant \sigma(w)$.*

Claim 5.9. *For a labeled vertex v, let $P(v) = L(x_1) y_1 \dots L(x_k) y_k L(v)$ be the decomposition of $P(v)$. If $w \in L(x_i)$, where $i \in \{1, \dots, k\}$, and $z \in P(v|y_i)$ then $(w, z) \notin E^\sharp$.*

Proof. Suppose, to the contrary, that $(w, z) \in E^\sharp$, where $w \in L(x_i)$ and $z \in P(v|y_i)$. Note that x_i is scanned, so, by Claims 5.6 and 5.5, w is scanned and $\sigma(w) \leqslant \sigma(x_i) = \lambda(y_i)$. Suppose that z is a transform; say $z = \tau(p, q, b)$. By Claim 5.7, $\sigma(b) \geqslant \lambda(y) \geqslant \sigma(w)$. However, since w is adjacent to z, $w \neq b$. Hence, $\sigma(b) > \sigma(w)$.

Note that w is adjacent to either p or q; say p. If $\lambda(p)$ is defined, then $\lambda(p) = \sigma(b) > \sigma(w)$, contradicting Claim 5.8. Therefore, we may assume that $\lambda(p)$ is not defined. If z is created by $\mathsf{Blossom}(r, u)$, then p must belong to $P(r)$, which contradicts Claim 5.3. Therefore, $B(z)$ is a degenerate blossom, and, hence, q is the mate of p. However, $\lambda(q)$ is defined and $\lambda(q) = \sigma(b) > \sigma(w)$, contradicting Claim 5.8. Therefore, z is not a transform.

If $\lambda(z)$ is defined then, by Claim 5.5, $\lambda(z) \geqslant \lambda(y_i) \geqslant \sigma(w)$, contradicting Claim 5.8. Thus, $\lambda(z)$ is not defined. Then, by Claim 5.8, $\lambda(\bar{z})$ is defined and $\lambda(\bar{z}) \leqslant \sigma(w) \leqslant \lambda(y_i)$. Therefore, by Claim 5.5, we see that $\bar{z} \notin P(v)$. Then, there exists a transform $z' = \tau(\bar{z}, q, b)$ in $P(v)$. Since z' is labeled after z, we have $z \in L(v)$. By Claim 5.7, $\lambda(\bar{z}) = \sigma(b) \geqslant \lambda(y_i)$. This contradiction completes the proof. $\square$

Claim 5.10. *Suppose $w \in P(v)$ and $z \in P(u)$ such that $\sigma(w) \leqslant \sigma(v)$, $\sigma(z) \leqslant \sigma(u)$, and either $P(v) \cap B^\bullet(z) = \emptyset$ or $P(u) \cap B^\bullet(w) = \emptyset$ when the algorithm starts* $\mathsf{Blossom}(v, u)$. *If w is adjacent to z in $G^\sharp$, then $w = v$ and $z = u$.*

Proof. If (w, z) is such an edge and $(w, z) \neq (u, v)$, then the algorithm should have performed either $\mathsf{Blossom}(w, z)$ or $\mathsf{Blossom}(z, w)$ before $\mathsf{Blossom}(v, u)$. $\square$

We now intend to show that a parity path P obtained in Step 0 of $\mathsf{Blossom}(v, u)$ is in fact an augmenting path. Since (AP1)–(AP4) are obvious from the algorithm description, it suffices to show that P satisfies (AP5). In order prove (AP5) we utilize Lemma 2.1.

Claim 5.11. *For a labeled vertex v, there exists a partition $(L(x_i) \cup \{y_i\}$: $i = 1, \ldots, h)$ of $L(v) - \{v\}$ such that, for $i = 1, \ldots, h$, either $\sigma(x_i) = \lambda(y_i)$ or the algorithm has applied* $\mathsf{Blossom}(y_i, x_i)$. *Moreover, there is an appropriate total order $\preccurlyeq$ among the components such that $i \prec j$ implies $(w, z) \notin E^\sharp$ for every $w \in L(x_i)$ and $z \in L(x_j) \cup \{y_j\}$.*

Proof. Suppose inductively that the statements hold before v gets labeled. If $B(v)$ is degenerate then $L(v) = v$, so the result is vacuous. Now suppose that v is labeled in $\mathsf{Blossom}(r, u)$. Note that $v \in P(u)$, $v \in P(r)$, or v is a transform. We assume that $v \in P(u)$; similar arguments work for the other cases. We also assume that $L(v) \neq P(v)$; the argument when $L(v) = P(v)$ is similar.

Let y_0 be defined so that $L(v) = P(v|y_0)$. By Claim 5.5, $\lambda(y_0) \leqslant \lambda(v)$. Therefore, again by Claim 5.5, $y_0 \notin P(u|v)$. Thus, $L(v) = P(r|y_0)\overline{P(u|v)}v$. Since y_0 was not labeled when we labeled r, we can decompose $P(r)$ as $P(x_0)y_0L(x_1)y_1 \ldots L(x_k)y_kL(r)$. Similarly, $P(u|v)$ can be decomposed as $P(u|v) = L(x_{k+1})y_{k+1} \ldots L(x_{\ell-1})y_{\ell-1}L(x_\ell)$, where $x_\ell = u$ and $k < \ell$. By the inductive assumption $L(r) - \{r\}$ can be partitioned into $(L(x_i) \cup \{y_i\}$: $i = \ell + 1, \ldots, h)$. Putting $y_\ell = r$, we obtain a partition $(L(x_i) \cup \{y_i\}$: $i = 1, \ldots, h)$ of $L(v) - \{v\}$ such that, for each $i = 1, \ldots, h$, the algorithm has applied $\mathsf{Blossom}(y_i, x_i)$ or $\sigma(x_i) = \lambda(y_i)$. In particular, $\sigma(x_i) = \lambda(y_i)$ holds for $1 \leqslant i < \ell$.

By the inductive assumption, there is a total order $\preccurlyeq$ among the components in $L(r) - \{r\}$ such that $i \prec j$ implies $(w, z) \notin E^\sharp$ for every $w \in L(x_i)$ and $z \in L(x_j) \cup \{y_j\}$.

We extend this to a total ordering on all of the components by defining $i \prec j$ if either:

(a) $1 \leqslant i < \ell \leqslant j \leqslant k$ or
(b) $\sigma(x_i) < \sigma(x_j)$ where $1 \leqslant i, j \leqslant \ell$.

Now consider any $w \in L(x_i)$ and $z \in L(x_j)$, where $i \prec j$. If $i,\ j > \ell$ then, by definition, $(w, z) \notin E^\sharp$. By Claims 5.9 and 5.10, if $i, j \leqslant \ell$, we have $(w, z) \notin E^\sharp$. Now suppose that $1 \leqslant i \leqslant \ell < j \leqslant h$. Thus, $z \in P(r) - \{r\}$. So, by Claims 5.9 and 5.10, $(w, v) \notin E^\sharp$, as required. $\square$

Claim 5.12. *When the algorithm scans a labeled vertex v, the following hold.*

(1) *If the algorithm applies* Step 3 *of* $\mathsf{Search}(F)$ *to an unlabeled vertex u adjacent to v, then $G^\sharp[L(v) \cup \{u\}]$ has a unique perfect matching.*
(2) *If the algorithm applies* $\mathsf{Blossom}(v, u)$ *in* Step 2 *of* $\mathsf{Search}(F)$, *then $G^\sharp[L(u) \cup \{v\}]$ has a unique perfect matching.*

Proof. Suppose inductively that both statements hold for previously scanned vertices.

In order to prove (1), we consider the partition $(L(x_i) \cup \{y_i\}\colon i = 1, \dots, h)$ of $L(v) - \{v\}$ with the total order $\preccurlyeq$ introduced in Claim 5.11. By the inductive assumption, each induced subgraph $G^\sharp[L(x_i) \cup \{y_i\}]$ has a unique perfect matching.

Put $y_0 = v$, $L_0 = \{u\}$, and $L_i = L(x_i)$ for $i = 1, \dots, h$. We extend the total order $\preccurlyeq$ by setting $0 \prec i$ for $i = 1, \dots, h$. Now $\lambda(u) = \sigma(v)$, and, for $w \in L(v) - \{v\}$, we have $\sigma(w) < \sigma(v)$, so, by Claim 5.8, $(w, u) \notin E^\sharp$. Thus, $i \prec j$ implies $(w, z) \notin E^\sharp$ for $w \in L_i$ and $z \in L_j \cup \{y_j\}$. Therefore, by Lemma 2.1, $G^\sharp[L(u) \cup \{v\}]$ has a unique perfect matching.

In order to prove (2) we consider the partition of $L(u) - \{u\}$ into $L(x_i) \cup \{y_i\}$ for $i = 1, \dots, h$. For each $w \in L(u) - \{u\}$, we have $\sigma(w) < \sigma(u)$. Therefore, by Claim 5.10, $(w, v) \notin E^\sharp$. Put $y_0 = u$, $L_0 = \{v\}$, and $L_i = L(x_i)$ for $i = 1, \dots, h$. Then, as above, we see that (2) follows from Lemma 2.1. $\square$

Lemma 5.2. *The parity path P obtained by* $\mathsf{Search}(F)$ *is an augmenting path.*

Proof. Suppose the algorithm returns $P = P(v)\overline{P(u)}$ at Step 0 in $\mathsf{Blossom}(v, u)$. It suffices to show that $G^\sharp[P]$ has a unique perfect matching. The search paths $P(v)$ and $P(u)$ can be decomposed as $P(v) = L(x_1)y_1 \dots L(x_k)y_k L(v)$ and $P(u) = L(x_{k+1})y_{k+1} \dots L(x_{\ell-1})y_{\ell-1}L(x_\ell)$, where $x_\ell = u$ and $k < \ell$. Let $(L(x_i) \cup \{y_i\}\colon i = \ell + 1, \dots, h)$ be the partition of $L(v) - \{v\}$ given by Claim 5.11.

Putting $y_\ell = v$, we see that the set of vertices in P is a disjoint union of $L(x_i) \cup \{y_i\}$ for $i = 1, \dots, h$. By Claim 5.12, each component $G^\sharp[L(x_i) \cup \{y_i\}]$ has a unique perfect matching. Similar to the proof of Claim 5.12, we may define a total order $\preccurlyeq$ among the components such that $i \prec j$ implies $(w, z) \notin E^\sharp$ for every $w \in L(x_i)$ and $z \in L(x_j) \cup \{y_j\}$. Therefore, $G^\sharp[P]$ has a unique perfect matching by Lemma 2.1. $\square$

5.1. Proof of the minimax theorem

We now show that, if the algorithm fails to find an augmenting path, F is an optimum feasible set. Recall that for a matrix A we denote by $A[X, Y]$ a submatrix whose rows and columns are indexed by X and Y, respectively. In addition, we denote $A[X, X]$ by $A[X]$.

Let T_B denote the set of the tips of a blossom B with a bud. Construct a graph $\Gamma_B = (T_B, C_B)$ with vertex set T_B and edge set $C_B = \{(p, q) \mid \tau(p, q, b) \in B\}$. Note that, before we create a transform $\tau(p, q, b)$ there was no blossom that contained both p and q. Therefore, the graph Γ_B is a tree, which we call the *tip tree* of B.

When the algorithm generates a blossom B' that includes a blossom B and does not share the bud, one of the tips of B, say x, gets labeled, but the other tips of B do not. The following fact will be used in the proof of the minimax theorem.

Claim 5.13. *A vertex $z \in V^\sharp$ that is adjacent in $G^\sharp$ to an unlabeled tip y of the non-maximal blossom B is adjacent to the labeled tip x of B or some transform in B.*

Proof. The column corresponding to a tip y of B is a linear combination of those columns for x and for the transforms along the path in the tip tree Γ_B between y and x. Hence, for $z \in V^\sharp$, if $A^\sharp_{zy} \neq 0$ then $A^\sharp_{zw} \neq 0$ for $w = x$ or for some transform w in B. That is, a vertex $z \in V^\sharp$ that is adjacent in $G^\sharp$ to an unlabeled tip y of B is also adjacent to the labeled tip x of B or some transform in B. $\square$

Henceforth, we suppose that the algorithm $\mathsf{Search}(F)$ has terminated at Step 1 without finding an augmenting path. We intend to show that F is optimal to the delta-parity problem. Recall that a vertex in a blossom has its mate in the same blossom. We consider a partition of $V^\sharp$ defined by

$R^\sharp_i$: a maximal blossom that is disjoint of source lines ($i = 1, \ldots, k$);

$W^\sharp_j$: a maximal blossom that includes a source line ($j = 1, \ldots, \ell$);

U: the set of unlabeled elements with unlabeled mates.

Put $R_i = R^\sharp_i \cap V$ for $i = 1, \ldots, k$ and $W_j = W^\sharp_j \cap V$ for $j = 1, \ldots, \ell$. Then $(W_1, \ldots, W_\ell, U, R_1, \ldots, R_k)$ is a partition of V. Each $W^\sharp_j$ includes exactly one source line, which implies $\ell = \delta_\Pi(F)$. Let b_i denote the bud of $R^\sharp_i$ and T_i the set of the tips of $R^\sharp_i$. We also denote $R = \bigcup_{i=1}^k R_i$ and $W = \bigcup_{j=1}^\ell W_j$.

Claim 5.14. *The partition $(W_1, \ldots, W_\ell, U, R_1, \ldots, R_k)$ of V satisfies*:

(1) $A_F[W_i, W_j] = 0$ *and* $A_F[W_j, W_i] = 0$ *for* $1 \leqslant i < j \leqslant \ell$;
(2) $A_F[U, W] = 0$ *and* $A_F[W, U] = 0$.

Proof. Recall that $A_F = A^\sharp[V]$. Since W_i has no tips, it follows from Claim 5.13 that a vertex in $V^\sharp$ adjacent in $G^\sharp$ to some vertex in W_i is adjacent to a labeled vertex in W_i.

If there is an edge in $G^\sharp$ between two labeled vertices, these two must be in the same blossom. Therefore, the maximal blossoms W_i for $i = 1, \dots, \ell$ are not adjacent to each other, which implies (1). If a vertex $u \in U$ is adjacent to a vertex in W, it is adjacent to a labeled vertex in W, and then the mate $\bar{u}$ must be labeled or u is in a blossom, which contradicts the definition of U. Thus, we obtain (2). $\square$

Claim 5.15. *Suppose $u \in V^\sharp \backslash R_i^\sharp$ is labeled or belongs to U. Then u is not adjacent in $G^\sharp$ to any labeled vertex in $R_i^\sharp$.*

Proof. If $u \in V^\sharp \backslash R_i^\sharp$ is labeled and adjacent to a labeled vertex $x \in R_i^\sharp$, then u and x must be in the same blossom, which is a contradiction. On the other hand, if u belongs to U and is adjacent to a labeled vertex in $R_i^\sharp$, then the mate $\bar{u}$ must be labeled or u is in a blossom, which contradicts the definition of U. $\square$

Claim 5.16. *Suppose $u \in V^\sharp \backslash R_i^\sharp$ is not adjacent to any labeled vertex in $R_i^\sharp$. Then $A^\sharp[\{u\}, R_i]$ is a multiple of $A^\sharp[\{b_i\}, R_i]$.*

Proof. It follows from Claim 5.13 that $A^\sharp_{uv} = 0$ for every $v \in R_i^\sharp \backslash T_i$. On the other hand, since the bud $b_i \in V^\sharp \backslash R_i^\sharp$ of $R_i^\sharp$ is also labeled, we have $A^\sharp_{b_i v} = 0$ for every $v \in R_i^\sharp \backslash T_i$ by the same reasoning.

Next we consider the parts corresponding to T_i. Let r be a tip of $R_i^\sharp$. For another tip $p \in T_i$, consider the next tip $q \in T_i$ along the path from p to r in the tip tree. Then the transform t generated with reference to p and q is not adjacent to u, thus, $A^\sharp_{ut} = 0$. By the definition of $A^\sharp_{ut}$, for the transform t, we have $A^\sharp_{up} = A^\sharp_{uq} A^\sharp_{b_i p} / A^\sharp_{b_i q}$. Repeating this, we obtain $A^\sharp_{up} = A^\sharp_{ur} A^\sharp_{b_i p} / A^\sharp_{b_i r}$ for every tip $p \in T_i$. Thus, $A^\sharp[\{u\}, R_i]$ is a multiple of $A^\sharp[\{b_i\}, R_i]$. $\square$

We now construct a skew-symmetric matrix A^+, whose row/column set is $V \cup Z$ with $Z = \{z_i \mid i = 1, \dots, k\}$. The matrix A^+ is defined by $A^+[V] = A_F$, $A^+[Z] = 0$, $A^+[\{z_i\}, R_i] = A_F[\{b_i\}, R_i]$ and $A^+[\{z_i\}, V \backslash R_i] = 0$ for $i = 1, \dots, k$. We omit here the definitions of submatrices determined by the skew-symmetry.

Recall that A^+ is said to be congruent to A° if $A^\circ = Q^t A^+ Q$ holds for some nonsingular Q. In particular, if A° is a skew-symmetric matrix that is obtained from A^+ by adding multiples of row/column i to other rows/columns, then A° is congruent to A^+.

Claim 5.17. *The skew-symmetric matrix A^+ is congruent to a skew-symmetric matrix A° such that*

(1) $A^\circ[W_i, W_j] = 0$ *and* $A^\circ[W_j, W_i] = 0$ *for* $1 \leqslant i < j \leqslant \ell$;

(2) $A^{\circ}[U, W] = 0$ *and* $A^{\circ}[W, U] = 0$;
(3) $A^{\circ}[R, W] = 0$ *and* $A^{\circ}[W, R] = 0$;
(4) $A^{\circ}[R, U] = 0$ *and* $A^{\circ}[U, R] = 0$;
(5) $A^{\circ}[R_i, R_j] = 0$ *and* $A^{\circ}[R_j, R_i] = 0$ *for* $1 \leqslant i < j \leqslant k$;
(6) $A^{\circ}[Z, V] = A^{+}[Z, V]$, $A^{\circ}[V, Z] = A^{+}[V, Z]$, *and* $A^{\circ}[Z] = A^{+}[Z] = 0$.

Proof. We will obtain A° from A^{+} by adding multiples of the rows/columns of A^{+} that are indexed by Z to the rows/columns of A^{+} that are indexed by V. Since $A^{+}[W \cup U, Z] = 0$, this transformation does not change $A^{+}[W \cup U]$ nor $A^{+}[Z, Z \cup V]$, which together with Claim 5.14 implies (1), (2) and (6).

By Claims 5.13, 5.15 and 5.16, for each $u \in W \cup U$, $A^{\sharp}[\{u\}, R_i]$ is a multiple of $A^{\sharp}[\{b_i\}, R_i] = A^{+}[\{z_i\}, R_i]$. Hence, it is easy to make $A^{+}[W \cup U, R_i]$ zero, by a congruence transformation adding a multiple of the row/column indexed by z_i, without affecting any other part of A^{+}. Thus, we have (3) and (4).

Now consider eliminating $A^{+}[R_i, R_j]$ for $i \neq j$. By Claims 5.13 and 5.15, each vertex $v \in R_i^{\sharp} \backslash T_i$ is not adjacent to any labeled vertex in R_j. Hence, by Claim 5.16, $A^{\sharp}[\{v\}, R_j]$ is a multiple of $A^{\sharp}[\{b_j\}, R_j] = A^{+}[\{z_j\}, R_j]$ for $v \in R_i^{\sharp} \backslash T_i$. For each $v \in R_i^{\sharp} \backslash T_i$, we have $A^{\sharp}_{v b_i} = 0$. Therefore, by Claims 5.13 and 5.15, $A^{\sharp}[\{v\}, R_j \cup \{b_i\}]$ is a multiple of $A^{+}[\{z_j\}, R_j \cup \{b_i\}]$. Now, let p be a tip of the blossom $R_i^{\sharp}$. For any other tip q of $R_i^{\sharp}$, the column $A^{\sharp}_{*q}$ is a linear combination of $A^{\sharp}_{*p}$ and the columns of $A^{\sharp}$ that are indexed by transforms in $R_i^{\sharp}$. However, as shown above, for each transform t in $R_i^{\sharp}$, $A^{\sharp}[\{t\}, R_j \cup \{b_i\}]$ is a multiple of $A^{+}[\{z_j\}, R_j \cup \{b_i\}]$. Hence, $A^{+}[\{q\}, R_j \cup \{b_i\}]$ is a linear combination of $A^{+}[\{z_j\}, R_j \cup \{b_i\}]$ and $A^{+}[\{p\}, R_j \cup \{b_i\}]$ for every $q \in T_i$. Therefore, we can apply a congruence transformation to A^{+} so that, for every $x \in R_i$, $A^{+}[\{x\}, R_j \cup \{b_i\}]$ becomes a multiple of $A^{+}[\{p\}, R_j \cup \{b_i\}]$ and the other parts of A^{+} remain unchanged. As a result of this congruence transformation, we obtain rank $A^{+}[R_i, R_j \cup \{b_i\}] = 1$. In particular, $A^{+}[R_i, \{y\}]$ becomes a multiple of $A^{+}[R_i, \{b_i\}] = A^{+}[R_i, \{z_i\}]$ for every $y \in R_j$. Thus, we can eliminate $A^{+}[R_i, R_j]$.

We may apply the above argument separately for different pairs of i and j to obtain (5). $\square$

Claim 5.18. *Let* A° *be the matrix found in the proof of Claim* 5.17. *Then* $\mathbf{M}(A^{\circ})/Z = \mathbf{M}(A^{+})/Z$. *Moreover* $\mathbf{M}(A^{+})/Z$ *is nonempty.*

Proof. Let X be a subset of V. As A° is obtained from A^{+} by adding multiples of the rows/columns indexed by Z to those indexed by V, $A^{\circ}[Z \cup X]$ is congruent to $A^{\circ}[Z \cup X]$. Therefore, $A^{\circ}[Z \cup X]$ is nonsingular if and only if $A^{+}[Z \cup X]$ is nonsingular. Therefore, $\mathbf{M}(A^{\circ})/Z = \mathbf{M}(A^{+})/Z$.

The columns of A^{+} that are indexed by Z are linearly independent. Let $Z \cup X$ be a maximal set that indexes a set of linearly independent columns of A^{+}. Then, $A^{+}[Z \cup X]$ is nonsingular, and, hence, X is a feasible set of $\mathbf{M}(A^{+})/Z$. $\square$

Since $\mathbf{M}(A^+)\backslash Z = \mathbf{M}(A_F)$ and $\mathbf{M}(A_F) = \mathbf{M} \triangle F$, we have $\mathrm{dist}(\mathbf{M} \triangle F, \mathbf{M}(A^\circ)/Z) \leqslant |Z|$, which implies $\mathrm{dist}(\mathbf{M}, \mathbf{M}^\circ) \leqslant |Z|$ for $\mathbf{M}^\circ = (\mathbf{M}(A^\circ)/Z) \triangle F$.

We now evaluate $\mathrm{odd}(\mathbf{M}^\circ)$. Denote $R_j^+ = R_j \cup \{z_j\}$ for $j = 1, \ldots, k$. Then Claim 5.17 implies that the partition $(W_1, \ldots, W_\ell, U, R_1^+, \ldots, R_k^+)$ gives a direct-sum decomposition of $\mathbf{M}(A^\circ)$ without odd components. Therefore, the partition $(W_1, \ldots, W_\ell, U, R_1, \ldots, R_k)$ determines a direct-sum decomposition of $\mathbf{M}(A^\circ)/Z$ consistent with Π. The components corresponding to $R_1, \ldots, R_k$ have feasible sets of odd cardinality but the others do not.

The same partition gives a direct-sum decomposition of $\mathbf{M}^\circ$. Since W_i is a union of lines and includes exactly one source line, $|F \cap W_i|$ is odd for each $i = 1, \ldots, k$. Therefore, W_i corresponds to an odd component of $\mathbf{M}^\circ$. On the other hand, R_j corresponds to an odd component of $\mathbf{M}^\circ$ because $|F \cap R_j|$ is even. Thus, all the components of $\mathbf{M}^\circ$ except for U are odd components, which implies $\mathrm{odd}(\mathbf{M}^\circ, \Pi) \geqslant \ell + k$. Combining this with $\mathrm{dist}(\mathbf{M}, \mathbf{M}^\circ) \leqslant |Z| = k$, we obtain $\mathrm{odd}(\mathbf{M}^\circ, \Pi) - \mathrm{dist}(\mathbf{M}, \mathbf{M}^\circ) \geqslant \ell$, which together with Lemma 3.3 completes the proof of Theorem 3.1.

5.2. Complexity

We conclude by discussing the time complexity of the augmenting path algorithm. Obviously from the construction of the tip trees, the number of transforms generated in $\mathsf{Search}(F)$ is at most $n = |V|$. Hence, the size of the matrix we deal with is $O(n)$, and the search procedure can be implemented to run in $O(n^2)$ time. The augmentation is in fact a pivoting, which requires $O(n^3)$ time. Since the whole algorithm performs at most $n/4$ augmentations, the total time complexity is $O(n^4)$.

6. Related work

The delta-parity problem is equivalent to computing the rank of mixed skew-symmetric matrices introduced in [16, Section 7.3]. A subsequent paper [10] exhibits an alternative, nonalgorithmic approach to show the minimax theorem via the rank of mixed skew-symmetric matrices.

Acknowledgments

We are grateful to André Bouchet and András Sebő for helpful comments, which resulted in a substantial improvement of the presentation in Section 5. The first author thanks Bill Jackson for illuminating discussions on this subject.

References

[1] A. Bouchet, Greedy algorithm and symmetric matroids, Math. Programming 38 (1987) 147–159.
[2] A. Bouchet, Representability of Δ-matroids, Combinatorics, Colloq. Math. Soc. János Bolyai 52 (1988) 167–182.

[3] A. Bouchet, Coverings and delta-coverings, in: E. Balas, J. Clausen (Eds.), Integer Programming and Combinatorial Optimization, Lecture Notes in Computer Science, Vol. 920, Springer, Berlin, 1995, pp. 228–243.
[4] A. Bouchet, W.H. Cunningham, Delta-matroids, jump systems, and bisubmodular polyhedra, SIAM J. Discrete Math. 8 (1995) 17–32.
[5] A. Bouchet, A. Duchamp, Representability of Δ-matroids over GF(2), Linear Algebra Appl. 146 (1991) 67–78.
[6] R. Chandrasekaran, S.N. Kabadi, Pseudomatroids, Discrete Math. 71 (1988) 205–217.
[7] A. Dress, T. Havel, Some combinatorial properties of discriminants in metric vector spaces, Adv. Math. 62 (1986) 285–312.
[8] J. Edmonds, Submodular functions matroids, and certain polyhedra, in: R. Guy, H. Hanani, N. Sauer, J. Schönheim (Eds.), Combinatorial Structures and Their Applications, Gordon and Breach, New York, 1970, pp. 69–87.
[9] H.N. Gabow, M. Stallmann, An augmenting path algorithm for linear matroid parity, Combinatorica 6 (1986) 123–150.
[10] J.F. Geelen, S. Iwata, Matroid matching via mixed skew-symmetric matrices, METR 2002–03, Univ. of Tokyo, 2002.
[11] P.M. Jensen, B. Korte, Complexity of matroid property algorithms, SIAM J. Comput. 11 (1982) 184–190.
[12] L. Lovász, The matroid matching problem, Algebraic Methods in Graph Theory, Colloq. Math. Soc. János Bolyai 25 (1978) 495–517.
[13] L. Lovász, Matroid matching and some applications, J. Combin. Theory Ser. B 28 (1980) 208–236.
[14] L. Lovász, Selecting independent lines from a family of lines in a space, Acta Sci. Math. 42 (1980) 121–131.
[15] L. Lovász, M.D. Plummer, Matching Theory, North-Holland, Amsterdam, 1986.
[16] K. Murota, Matrices and Matroids for Systems Analysis, Springer, Berlin, 2000.
[17] J.B. Orlin, J.H. Vande Vate, Solving the linear matroid parity problem as a sequence of matroid intersection problems, Math. Programming 47 (1990) 81–106.
[18] J.G. Oxley, Matroid Theory, Oxford University Press, Oxford, 1992.
[19] A.W. Tucker, A combinatorial equivalence of matrices, Proc. Sympos. Appl. Math. 10 (1960) 129–140.
[20] D.J.A. Welsh, Matroid Theory, Academic Press, London, 1976.