The vector partition problem for convex objective functions

Shmuel Onn * Leonard J. Schulman †

Abstract

The *partition problem* concerns the partitioning of a given set of n vectors in d-space into p parts so as to maximize an objective function which is convex on the sum of vectors in each part. The problem has broad expressive power and captures NP-hard problems even if either p or d is fixed. In this article we show that when both p, d are fixed, the problem is solvable in strongly polynomial time using $O(n^{d(p-1)-1})$ arithmetic operations. This improves upon the previously known bound of $O(n^{dp^2})$. Our method is based on the introduction of the *signing zonotope* of a set of points in space. We study this object, which is of interest in its own right, and show that it is a refinement of the so called *partition polytope* of the same set of points.

Mathematics Subject Classifications: 05A, 15A, 51M, 52A, 52B, 52C, 68Q, 68R, 68U, 90B, 90C

Keywords: partition, cluster, optimization, separation, convex, polytope, zonotope, signing, vertex enumeration, polynomial time, combinatorial optimization

1 Introduction

The *partition problem* concerns the partitioning of a set $A = \{a_1, \ldots, a_n\}$ of n vectors in d-space into p parts so as to maximize an objective function which is convex on the sum of vectors in each part. Each ordered partition $\pi = (\pi_1, \ldots, \pi_p)$ of A is associated with the $d \times p$ matrix $A^\pi := \left[\sum_{a\in\pi_1} a, \ldots, \sum_{a\in\pi_p} a\right]$, whose jth column represents the total value of vectors assigned to the jth part. The problem is to find a partition π which maximizes an objective function f given by $f(\pi) = c(A^\pi)$, where c is a real convex functional on $\mathbb{R}^{d\times p}$. This class of problems has applications in diverse fields that include circuit layout, clustering, inventory, scheduling and reliability - see Barnes, Hoffman and Rothblum (1992), Granot and Rothblum (1991), Hwang, Onn and Rothblum (2000), and references therein - as well as important recent applications to symbolic computation (Onn and Sturmfels (1999)). In its full generality, the partition problem instantly captures NP-hard problems hence is presumably intractable (Hwang, Onn and Rothblum (1999)).

Since the problem concerns the maximization of a convex function, it can be reduced to the problem of maximizing the same objective over the *p-partition polytope* $\mathcal{P}^p_A$ of A defined to be the

*Research supported in part by a grant from the Israel Science Foundation, by a VPR grant at the Technion, and by the Fund for the Promotion of Research at the Technion.

†Supported in part by NSF CAREER grant CCR-98761722 and the Charles Lee Powell Foundation.

convex hull in $\mathbb{R}^{d\times p}$ of all (exponentially many) matrices A^π corresponding to p-partitions. As there will always be an optimal solution which is a vertex of this polytope, the partition problem can be solved by picking the best among the *extremal* partitions of A - those π with A^π a vertex of $\mathcal{P}^p_A$. If the convex functional c is presented by an evaluation oracle, it may be necessary in worst case to query the oracle on each and every one of these extremal partitions. Therefore, the complexity of the partition problem is intimately related to the number of extremal partitions.

As the number of partitions π and matrices A^π is exponential in n even with both p, d fixed, the filtration of extremal out of all partitions by direct enumeration is prohibitive. Barnes, Hoffman and Rothblum (1992) showed that any extremal partition is *separable*, that is, any two distinct parts can be separated by a hyperplane. This condition was exploited by Hwang, Onn and Rothblum (1999) who showed that for any fixed p, d, the number of separable and hence extremal partitions is $O(n^{d\binom{p}{2}})$, and all separable and hence extremal partitions can be enumerated and the partition problem solved in time $O(n^{dp^2})$. These results apply for the more general *shaped partition problem*, where partitions are restricted to be those π whose *shape* $(|\pi_1|, \ldots, |\pi_p|)$ lies in a prescribed (but arbitrary) set of shapes of n. The solution of the partition problem by this method can not be significantly accelerated: Alon and Onn (1999) showed that, for every fixed $p \geq 2$ and $d \geq 3$, the number of separable partitions is in fact $\Theta(n^{d\binom{p}{2}})$.

In this article we provide a more efficient resolution of the (unrestricted) partition problem by taking a different approach and introducing the *p-signing zonotope* $\mathcal{Z}^p_A$ of A defined to be the convex hull in $\mathbb{R}^{d\times p}$ of all matrices A_σ corresponding to *p-signings* σ of A; the formal definitions are provided in the next section, where we study this object and establish Theorem 2.4 which implies at once the following result, demonstrating a many-to-one mapping from the set of vertices of $\mathcal{Z}^p_A$ onto the set of vertices of $\mathcal{P}^p_A$:

Theorem 1.1 *For every vertex A_σ of the signing zonotope $\mathcal{Z}^p_A$ there exists a vertex A^π of the partition polytope $\mathcal{P}^p_A$ such that the cone of linear functionals uniquely maximized over $\mathcal{Z}^p_A$ at A_σ is contained in the cone of linear functionals uniquely maximized over $\mathcal{P}^p_A$ at A^π.*

In fact, as shown in Example 2.5, the number of extremal signings (those σ with A_σ a vertex of $\mathcal{Z}^p_A$) corresponding to each extremal partition is typically exponential. Nevertheless, there are far fewer extremal signings than separable partitions and so, combining the many-to-one correspondence implied by Theorem 1.1 with available algorithmic and extremal results on zonotopes, we are able in Section 3 to conclude the following improved combinatorial and algorithmic bounds for every fixed d and $p \geq 2$ (see Section 3 for the precise statements):

Corollary 3.1: The maximum number of extremal p-partitions of any set of n nonzero points in $\mathbb{R}^d$ satisfies $e_{p,d}(n) = O(n^{d(p-1)-1})$.

Corollary 3.2: All extremal p-partitions are enumerable in strongly polynomial time using $O(n^{d(p-1)-1})$ arithmetic operations.

Corollary 3.3: The p-partition problem with an oracle presented c can be solved in strongly polynomial time using $O(n^{d(p-1)-1})$ arithmetic operations and queries.

We conclude our article with Corollary 3.4 which further elaborates on the relation between $\mathcal{P}^p_A$

and $\mathcal{Z}_A^p$, showing that, in a precise sense, the later is a refinement of the former.

Recent results of Aviran, Lev-Tov, Onn and Rothblum (to appear) which study partition polytopes for $p = 2$ or $d = 2$ give $e_{2,d} = \Theta(n^{d-1})$ for all $d \geq 2$ and $e_{p,2} = \Theta(n^p)$ for all $p \geq 3$. Thus, the bound of Corollary 3.1 is tight for $p = 2$ but somewhat loose for $d = 2$. However, already for $p = d = 3$ the precise rate of growth is as yet unknown. Several interesting questions remain. How accurate are the bounds of Corollary 3.1 on $e_{p,d}(n)$ for $p, d \geq 3$? Are they tight for the maximum number of extremal p-signings of any set of n nonzero points in d-space (which by Theorem 1.1 is greater than or equal to the maximum number $e_{p,d}(n)$ of extremal p-partitions)? Could the extremal partitions be enumerated and the partition problem solved faster?

2 Extremal Partitions and Extremal Signings

We start with some basic notation and conventions. Throughout, e_i stands for the ith standard unit vector in Euclidean real space. The outer product of an ordered pair u, v of vectors is the matrix $u \otimes v$ whose (i, j)th entry is $u_i v_j$. The inner product of two matrices U, V of the same dimensions is $\langle U, V \rangle := \sum_{i,j} U_{i,j} \cdot V_{i,j}$. Slightly modifying standard terminology, we call a (possibly empty) subset $\mathcal{C}$ of real space a *cone* provided $\mathcal{C} \cup \{0\}$ is closed under finite nonnegative linear combinations (i.e. is a cone in the standard sense). We make the convention that a sum over an empty set of vectors (matrices) is the zero vector (matrix) of dimension which is clear from the context.

A *p-partition* of a set A of points in $\mathbb{R}^d$ is an ordered collection $\pi = (\pi_1, \ldots, \pi_p)$ of p pairwise disjoint (possibly empty) sets whose union is A. We also interpret π as the function from A to $\{1, \ldots, p\}$ with $\pi(a)$ being the index for which $a \in \pi_{\pi(a)}$. With each p-partition π of A we associate the $d \times p$ matrix

$$A^\pi \quad := \quad \sum_{a \in A} a \otimes e_{\pi(a)} \quad = \quad \left[\sum_{a \in \pi_1} a, \ldots, \sum_{a \in \pi_p} a \right] \quad \in \quad \mathbb{R}^{d \times p} \quad .$$

We are concerned with the following combinatorial optimization problem (with the reals $\mathbb{R}$ replaced by the rationals $\mathbb{Q}$ when the Turing computation model is considered).

Partition Problem. Given positive integers p, d, n, a set A of n points in $\mathbb{R}^d$, and a convex functional $c : \mathbb{R}^{d \times p} \longrightarrow \mathbb{R}$, find a p-partition π^* attaining the maximum objective value,

$$c(A^{\pi^*}) \quad = \quad \max \{ c(A^\pi) \ : \ \pi \text{ is a } p\text{-partition of } A \} \quad .$$

Since c is convex, there will always be an optimal partition π which is *extremal*, that is, whose matrix A^π is a vertex of the convex hull of all matrices of p-partitions, defined as follows.

Partition Polytope. The *p-partition polytope* of a set A of n points in $\mathbb{R}^d$ is defined to be the convex hull of all p^n matrices of p-partitions of A,

$$\mathcal{P}_A^p \quad := \quad \text{conv} \{ A^\pi \ : \ \pi \text{ is a } p\text{-partition of } A \} \quad \subset \quad \mathbb{R}^{d \times p} \quad .$$

While $\mathcal{P}^p_A$ is defined as the convex hull of exponentially many points, it was shown by Hwang, Onn and Rothblum (1999) (for the more general class of *shaped partition polytopes*), using *separable partitions*, that for any fixed d, p, the number of extremal partitions is polynomial in n and they can be enumerated in polynomial time. Here we take a different approach, based on a new object - the *signing zonotope* - which leads to improved bounds and to a more efficient enumeration procedure for extremal partitions . We now proceed to define this object.

A *p-signing* of a set A of n points in $\mathbb{R}^d$ is an $n\binom{p}{2}$-tuple $\sigma = (\sigma^a_{r,s})$ with $\sigma^a_{r,s} \in \{-1, 1\}$ for all $a \in A$ and $1 \le r < s \le p$. We extend the domain of σ to all pairs $1 \le r \ne s \le p$ by antisymmetry $\sigma^a_{s,r} := -\sigma^a_{r,s}$. With each p-signing σ of A we associate the $d \times p$ matrix

$$A_\sigma \quad := \quad \sum_{a\in A} \sum_{1\le r<s\le p} \sigma^a_{r,s} \cdot a \otimes (e_r - e_s) \quad \in \quad \mathbb{R}^{d\times p} \quad .$$

As with partitions, a crucial role will be played by signings σ which are *extremal*, that is, whose matrix A_σ is a vertex of the convex hull of all matrices of p-signings, defined as follows.

Signing Zonotope. The *p-signing zonotope* of a set A of n points in $\mathbb{R}^d$ is defined to be the convex hull of all $2^{n\binom{p}{2}}$ matrices of p-signings of A,

$$\mathcal{Z}^p_A \quad := \quad \mathrm{conv}\ \{\, A_\sigma \ :\ \sigma \text{ is a } p\text{-signing of } A \,\} \quad \subset \quad \mathbb{R}^{d\times p} \quad .$$

It is indeed a zonotope, being the following Minkowski sum of $n\binom{p}{2}$ line segments,

$$\mathcal{Z}^p_A \quad = \quad \sum_{a\in A} \sum_{1\le r<s\le p} [-1, 1] \cdot a \otimes (e_r - e_s) \quad \subset \quad \mathbb{R}^{d\times p} \quad .$$

Define the *cone* $\mathcal{C}^\pi$ of a p-partition π of A to be the normal cone of $\mathcal{P}^p_A$ at A^π, consisting of all $d \times p$ matrices C for which the linear functional $\langle C, \cdot\rangle$ is uniquely maximized over $\mathcal{P}^p_A$ at A^π. Thus, $\mathcal{C}^\pi$ is an open cone (see first paragraph of this section) which is nonempty if and only if π is extremal. Analogously, define the *cone* $\mathcal{C}_\sigma$ of a p-signing σ of A to be the normal cone of $\mathcal{Z}^p_A$ at A_σ, consisting of all $d \times p$ matrices C for which the linear functional $\langle C, \cdot\rangle$ is uniquely maximized over $\mathcal{Z}^p_A$ at A_σ. Again, $\mathcal{C}_\sigma$ is an open cone which is nonempty if and only if σ is extremal.

The following two lemmas provide characterizations of the above cones and hence, in particular, of the extremal partitions and extremal signings. To avoid trivial but tedious complications, we shall assume that the set A does not contain the origin 0, which is clearly irrelevant for the partition problem. The first lemma, which concerns the partition polytope, can be found also in Aviran, Lev-Tov, Onn and Rothblum (to appear); we state and prove it here in a form more suitable for the sequel.

Lemma 2.1 *The cone $\mathcal{C}^\pi$ of a p-partition π of a set A of nonzero points in $\mathbb{R}^d$ is the set of all $d \times p$ matrices C with $\langle C, a \otimes (e_{\pi(a)} - e_k)\rangle$ positive for all $a \in A$ and $k \ne \pi(a)$.*

Proof. Suppose first $C \in \mathcal{C}^\pi$ hence $\langle C, \cdot\rangle$ is uniquely maximized over $\mathcal{P}^p_A$ at A^π. Consider any $a \in A$ and $k \ne \pi(a)$, and let $\bar\pi$ be the partition obtained from π by moving a to π_k. Then

$A^\pi - A^{\bar\pi} = a \otimes (e_{\pi(a)} - e_k)$ hence

$$\langle C, a \otimes (e_{\pi(a)} - e_k)\rangle \quad = \quad \langle C, A^\pi - A^{\bar\pi}\rangle \quad = \quad \langle C, A^\pi\rangle - \langle C, A^{\bar\pi}\rangle \quad > \quad 0 \quad .$$

Conversely, suppose C is a matrix with $\langle C, a \otimes (e_{\pi(a)} - e_k)\rangle$ positive for all $a \in A$ and $k \neq \pi(a)$. Consider any partition $\bar\pi$ different from π. Then for all $a \in A$ we have the inequality $\langle C, a \otimes (e_{\pi(a)} - e_{\bar\pi(a)})\rangle \geq 0$, being strict if $\pi(a) \neq \bar\pi(a)$. Since $\bar\pi \neq \pi$, the inequality is indeed strict for some $a \in A$, hence

$$\langle C, A^\pi\rangle - \langle C, A^{\bar\pi}\rangle \quad = \quad \sum_{a \in A} \left\langle C, a \otimes (e_{\pi(a)} - e_{\bar\pi(a)})\right\rangle \quad > \quad 0 \quad .$$

It follows that $\langle C, \cdot\rangle$ is uniquely maximized over $\mathcal{P}_A^p$ at A^π hence $C \in \mathcal{C}^\pi$. □

Lemma 2.1 also implies that, if π and $\bar\pi$ are two distinct extremal p-partitions of A then $A^\pi \neq A^{\bar\pi}$, and so extremal partitions stand in bijection with vertices of $\mathcal{P}_A^p$. Indeed if $\pi \neq \bar\pi$ are extremal with $A^\pi = A^{\bar\pi}$ then, taking any $C \in \mathcal{C}^\pi = \mathcal{C}^{\bar\pi}$ and $a \in A$ with $\pi(a) \neq \bar\pi(a)$, we must have by the lemma that both $\langle C, a \otimes (e_{\pi(a)} - e_{\bar\pi(a)})\rangle$ and $\langle C, a \otimes (e_{\bar\pi(a)} - e_{\pi(a)})\rangle$ must be positive, which is impossible. Note however that, for nonextremal partitions, it may happen that exponentially many π give the same matrix A^π.

The second lemma is the analog of Lemma 2.1 concerning the signing zonotope.

Lemma 2.2 *The cone $\mathcal{C}_\sigma$ of a p-signing σ of a set A of nonzero points in $\mathbb{R}^d$ is the set of all $d \times p$ matrices C with $\langle C, \sigma_{r,s}^a \cdot a \otimes (e_r - e_s)\rangle$ positive for all $a \in A$ and $1 \leq r < s \leq p$.*

Proof. Suppose first $C \in \mathcal{C}_\sigma$ hence $\langle C, \cdot\rangle$ is uniquely maximized over $\mathcal{Z}_A^p$ at A_σ. Consider any $a \in A$ and $1 \leq r < s \leq p$, and let $\bar\sigma$ be the signing obtained from σ by flipping the sign of $\sigma_{r,s}^a$. Then $A_\sigma - A_{\bar\sigma} = 2\sigma_{r,s}^a \cdot a \otimes (e_r - e_s)$ hence

$$\langle C, \sigma_{r,s}^a \cdot a \otimes (e_r - e_s)\rangle \quad = \quad \langle C, \frac{1}{2}(A_\sigma - A_{\bar\sigma})\rangle \quad = \quad \frac{1}{2}(\langle C, A_\sigma\rangle - \langle C, A_{\bar\sigma}\rangle) \quad > \quad 0 \quad .$$

Conversely, suppose C is a matrix with $\langle C, \sigma_{r,s}^a \cdot a \otimes (e_r - e_s)\rangle$ positive for all $a \in A$ and $1 \leq r < s \leq p$. Consider any signing $\bar\sigma$ different from σ. Then for all $a \in A$ and $1 \leq r < s \leq p$ we have the inequality $\langle C, (\sigma_{r,s}^a - \bar\sigma_{r,s}^a) \cdot a \otimes (e_r - e_s)\rangle \geq 0$, being strict if $\sigma_{r,s}^a \neq \bar\sigma_{r,s}^a$. Since $\sigma \neq \bar\sigma$, the inequality is indeed strict for some $a \in A$ and $1 \leq r < s \leq p$, hence

$$\langle C, A_\sigma\rangle - \langle C, A_{\bar\sigma}\rangle \quad = \quad \sum_{a \in A} \sum_{1 \leq r < s \leq p} \left\langle C, (\sigma_{r,s}^a - \bar\sigma_{r,s}^a) \cdot a \otimes (e_r - e_s)\right\rangle \quad > \quad 0 \quad .$$

It follows that $\langle C, \cdot\rangle$ is uniquely maximized over $\mathcal{Z}_A^p$ at A_σ hence $C \in \mathcal{C}_\sigma$. □

As for partitions, Lemma 2.2 also implies that, if σ and $\bar\sigma$ are two distinct extremal p-signings of A then $A_\sigma \neq A_{\bar\sigma}$, and so extremal signings stand in bijection with vertices of $\mathcal{Z}_A^p$. Indeed if $\sigma \neq \bar\sigma$ are extremal with $A_\sigma = A_{\bar\sigma}$ then, taking any $C \in \mathcal{C}_\sigma = \mathcal{C}_{\bar\sigma}$, and taking a point $a \in A$ and indices $1 \leq r < s \leq p$ with $\sigma^a_{r,s} \neq \bar{\sigma}^a_{r,s}$, we must have by the lemma that both $\langle C, \sigma^a_{r,s} \cdot a \otimes (e_r - e_s)\rangle$ and $\langle C, \bar{\sigma}^a_{r,s} \cdot a \otimes (e_r - e_s)\rangle = -\langle C, \sigma^a_{r,s} \cdot a \otimes (e_r - e_s)\rangle$ are positive, which is impossible. Note however that, for nonextremal signings, it may happen that exponentially many σ give the same matrix A_σ.

We proceed to determine a relation between extremal partitions and extremal signings.

Proposition 2.3 *If the p-signing σ of a set A of nonzero points is extremal then for every $a \in A$ there is a unique index $1 \leq \pi(a) \leq p$ satisfying $\sigma^a_{\pi(a),k} = 1$ for all $k \neq \pi(a)$.*

Proof. Let C be any matrix in the nonempty cone $\mathcal{C}_\sigma$ of the extremal signing σ. Consider any $a \in A$. By Lemma 2.2, for all $1 \leq r \neq s \leq p$, we have $\langle C, a \otimes e_r\rangle > \langle C, a \otimes e_s\rangle$ if and only if $\sigma^a_{r,s} = 1$. Thus, the values $\langle C, a \otimes e_1\rangle, \ldots, \langle C, a \otimes e_p\rangle$ are distinct. Let $\pi(a)$ be the unique index attaining maximum value $\langle C, a \otimes e_{\pi(a)}\rangle$. Then for all $k \neq \pi(a)$ we have $\langle C, a \otimes e_{\pi(a)}\rangle > \langle C, a \otimes e_k\rangle$ hence $\sigma^a_{\pi(a),k} = 1$, and so $\pi(a)$ is the desired unique index. □

Proposition 2.3 enables to associate with each extremal p-signing σ of A a p-partition $\pi := \pi[\sigma]$ of A via the interpretation of π as a function from A to $\{1, \ldots, p\}$ as follows: for each $a \in A$ let $1 \leq \pi(a) \leq p$ be the unique index satisfying $\sigma^a_{\pi(a),k} = 1$ for all $k \neq \pi(a)$. If σ is not extremal then $\pi[\sigma]$ is undefined. The following statement implies at once Theorem 1.1 stated in Section 1.

Theorem 2.4 *A p-partition π of a set A of nonzero points in $\mathbb{R}^d$ is extremal if and only if $\pi = \pi[\sigma]$ for some extremal p-signing σ of A. Moreover, the closure of the cone $\mathcal{C}^\pi$ of any p-partition π of A is the union of closures of all cones $\mathcal{C}_\sigma$ of p-signings σ of A with $\pi = \pi[\sigma]$.*

Proof. First, consider any extremal p-signing σ of A and let $\pi = \pi[\sigma]$ be its associated p-partition. We show that the corresponding cones satisfy $\mathcal{C}_\sigma \subseteq \mathcal{C}^\pi$ hence, in particular, $\mathcal{C}^\pi$ is nonempty and π is extremal. Consider any $C \in \mathcal{C}_\sigma$. For all $a \in A$ and $k \neq \pi(a)$ we have, by the definition of $\pi = \pi[\sigma]$, that $\sigma^a_{\pi(a),k} = 1$ and hence

$$\langle C\, ,\, a \otimes (e_{\pi(a)} - e_k)\rangle \quad = \quad \langle C\, ,\, \sigma^a_{\pi(a),k} \cdot a \otimes (e_{\pi(a)} - e_k)\rangle \quad .$$

Since $C \in \mathcal{C}_\sigma$ Lemma 2.2 implies that the right hand side is positive, hence so is the left hand side. Since this is true for all $a \in A$ and $k \neq \pi(a)$, Lemma 2.1 implies that $C \in \mathcal{C}^\pi$.

Next, let π be any extremal p-partition. Suppose σ is a p-signing for which the closure of $\mathcal{C}_\sigma$ intersects $\mathcal{C}^\pi$. Since both $\mathcal{C}^\pi$ and $\mathcal{C}_\sigma$ are open, $\mathcal{C}_\sigma$ intersects $\mathcal{C}^\pi$ as well. Since the cones of any two distinct extremal p-partitions are disjoint, and, as just shown in the first part of the proof above, $\mathcal{C}_\sigma \subseteq \mathcal{C}^{\pi[\sigma]}$, it follows that $\pi = \pi[\sigma]$. Now, the union over the finitely many p-signings σ of the closures of the cones $\mathcal{C}_\sigma$ is the entire space, hence contains $\mathcal{C}^\pi$. Since we have shown that whenever the closure of $\mathcal{C}_\sigma$ intersects $\mathcal{C}^\pi$ then $\pi = \pi[\sigma]$ and $\mathcal{C}_\sigma \subseteq \mathcal{C}^\pi$, it follows that, as claimed, the closure of $\mathcal{C}^\pi$ is the union of closures of all cones $\mathcal{C}_\sigma$ of p-signings σ of A with $\pi = \pi[\sigma]$. □

We finish this section by demonstrating, for every n, a set of n points in $\mathbb{R}^n$ each extremal p-partition π of which equals $\pi[\sigma]$ for exponentially many extremal p-signings σ of this set.

Example 2.5 Let p, n be positive integers and let $A := \{e_1, \ldots, e_n\}$ be the set of unit vectors in $\mathbb{R}^n$. Then the matrices A^π of p-partitions of A are precisely all $\{0,1\}$-matrices with one 1 per row hence all are vertices of $\mathcal{P}_A^p$. Thus, all p^n-many p-partitions of A are extremal. Next, the matrices A_σ of extremal p-signings of A are precisely all matrices with each row a permutation of $[\,p-1, p-3, \ldots, 3-p, 1-p\,]$. Thus, $p!^n$-many p-signings of A are extremal. Interestingly, here $\mathcal{Z}_A^p$ is equivalent to the n-fold product of the *permutohedron* of order p. The matrix A^π of an extremal $\pi = \pi[\sigma]$ is obtained from the matrix A_σ of σ by replacing the maximal entry $(p-1)$ in each row by 1 and replacing all other entries by 0. Therefore, every extremal partition π satisfies $\pi = \pi[\sigma]$ for precisely $(p-1)!^n$-many extremal signings σ.

Specifically, consider the case $n = p = 3$, $A = \{e_1, e_2, e_3\} \subset \mathbb{R}^3$. Then $\mathcal{P}_A^p \subset \mathbb{R}^{3\times 3}$ has $3^3 = 27$ vertices while $\mathcal{Z}_A^p \subseteq \mathbb{R}^{3\times 3}$ has $3!^3 = 216$ vertices. Let $\pi := (\{e_1\}, \{e_2\}, \{e_3\})$ be the extremal 3-partition with $A^\pi = I_3$ the 3×3-identity. Encoding each 3-signing σ by the 9-tuple

$$(\sigma_{1,2}^{e_1}\ \sigma_{1,3}^{e_1}\ \sigma_{2,3}^{e_1}\ \sigma_{1,2}^{e_2}\ \sigma_{1,3}^{e_2}\ \sigma_{2,3}^{e_2}\ \sigma_{1,2}^{e_3}\ \sigma_{1,3}^{e_3}\ \sigma_{2,3}^{e_3}) \quad ,$$

the $(3-1)!^3 = 8$ extremal 3-signings σ with $\pi = \pi[\sigma]$, together with their matrices A_σ, are

$$\underset{(+++-+++--)}{\begin{bmatrix} 2 & 0 & -2 \\ 0 & 2 & -2 \\ 0 & -2 & 2 \end{bmatrix}} \quad \underset{(+++-++---)}{\begin{bmatrix} 2 & 0 & -2 \\ 0 & 2 & -2 \\ -2 & 0 & 2 \end{bmatrix}} \quad \underset{(+++--++--)}{\begin{bmatrix} 2 & 0 & -2 \\ -2 & 2 & 0 \\ 0 & -2 & 2 \end{bmatrix}} \quad \underset{(+++--+---)}{\begin{bmatrix} 2 & 0 & -2 \\ -2 & 2 & 0 \\ -2 & 0 & 2 \end{bmatrix}}$$

$$\underset{(++--+++--)}{\begin{bmatrix} 2 & -2 & 0 \\ 0 & 2 & -2 \\ 0 & -2 & 2 \end{bmatrix}} \quad \underset{(++--++---)}{\begin{bmatrix} 2 & -2 & 0 \\ 0 & 2 & -2 \\ -2 & 0 & 2 \end{bmatrix}} \quad \underset{(++---++--)}{\begin{bmatrix} 2 & -2 & 0 \\ -2 & 2 & 0 \\ 0 & -2 & 2 \end{bmatrix}} \quad \underset{(++---+---)}{\begin{bmatrix} 2 & -2 & 0 \\ -2 & 2 & 0 \\ -2 & 0 & 2 \end{bmatrix}} \quad .$$

The closure of the normal cone $\mathcal{C}^\pi$ of $\mathcal{P}_A^p$ at its vertex $I_3 = A^\pi$ is the union of the closures of the normal cones $\mathcal{C}_\sigma$ of $\mathcal{Z}_A^p$ at its eight vertices given by the matrices A_σ above.

3 Consequences

Combining Theorem 2.4 with available algorithmic and extremal results on zonotopes, we obtain the following consequences concerning extremal partitions and the partition problem we started with. The case $p = 1$ being straightforward, we restrict attention to $p \geq 2$.

Corollary 3.1 *For every fixed d and $p \geq 2$, the maximum number of extremal p-partitions of any set of n nonzero points in $\mathbb{R}^d$ satisfies $e_{p,d}(n) = O(n^{d(p-1)-1})$.*

Proof. It is known (see e.g. Gritzmann and Sturmfels (1993) and references therein) that the number of vertices of any D-dimensional zonotope which is the Minkowski sum of N line segments is at most $2\sum_{k=0}^{D-1}\binom{N-1}{k}$. Consider any set A of n nonzero points in d-space. Then the p-signing zonotope $\mathcal{Z}_A^p$ of A is the sum of $N := n\binom{p}{2}$ line segments. As is easy to see, each matrix A_σ has row sum zero, hence $\mathcal{Z}_A^p$ lives in the subspace of $d \times p$ matrices with this property, which has dimension $D := d(p-1)$. By Theorem 2.4, the number of extremal p-partitions of A is less than or equal to the number of extremal p-signings of A which - see discussion following Lemma 2.2 - equals the number of vertices of $\mathcal{Z}_A^p$. For fixed p, d we therefore obtain, as claimed, the upper bound

$$e_{p,d}(n) \quad \leq 2\sum_{k=0}^{D-1}\binom{N-1}{k} \quad = \quad 2\sum_{k=0}^{d(p-1)-1}\binom{n\binom{p}{2}-1}{k} \quad = \quad O\left(n^{d(p-1)-1}\right) \quad . \quad \square$$

Aviran, Lev-Tov, Onn and Rothblum (to appear), which study partition polytopes with $p = 2$ or $d = 2$, show that $e_{2,d}(n) = \Theta(n^{d-1})$ for every $d \geq 2$ and $e_{p,2}(n) = \Theta(n^p)$ for every $p \geq 3$. Thus, the bound $e_{2,d}(n) = O(n^{d-1})$ of Corollary 3.1 is tight for all d whereas the bound $e_{p,2}(n) = O(n^{2p-3})$ is tight for no $p \geq 4$. How accurate are the bounds of Corollary 3.1 on $e_{p,d}(n)$ for $p, d \geq 3$? Are they tight for the maximum number of extremal p-signings of any set of n nonzero points in d-space (which by Theorem 2.4 is greater than or equal to the maximum number $e_{p,d}(n)$ of extremal p-partitions) ?

The next two algorithmic corollaries concern both the real and Turing computation models: an algorithm is *strongly polynomial time* if it uses a number of real arithmetic operations polynomially bounded in n, and for rational data runs in time polynomial in the bit size of the input.

Corollary 3.2 *For every fixed d and $p \geq 2$, all extremal p-partitions of any set of n nonzero points in d-space are enumerable in strongly polynomial time using $O(n^{d(p-1)-1})$ arithmetic operations.*

Proof. It is known (see Gritzmann and Sturmfels (1993) and references therein) that for any fixed D there is a strongly polynomial time algorithm which, given any zonotope in $\mathbb{R}^D$ presented as the Minkowski sum $\mathcal{Z} := \sum_{k=1}^{N}[-1,1]\cdot v_k$ of N line segments, enumerates all sign vectors $\sigma \in \{-1,1\}^N$ corresponding to vertices $\sum_{k=1}^{N}\sigma_k v_k$ of $\mathcal{Z}$, using $O(N^{d-1})$ arithmetic operations.

Fix p and d and set $D := d(p-1)$. Let A be any given set of n nonzero points in d-space. Let $N := n\binom{p}{2}$. For each $a \in A$ and $1 \leq r < s \leq p$ let $V_{r,s}^a$ be the $d \times (p-1)$ matrix obtained by omitting the last column of the $d \times p$ matrix $a \otimes (e_r - e_s)$. Let

$$\mathcal{Z} \quad := \quad \sum_{a\in A}\sum_{1\leq r<s\leq p}[-1,1]\cdot V_{r,s}^a \quad \subset \quad \mathbb{R}^{d\times(p-1)}$$

be the resulting sum of N line segments in $\mathbb{R}^D$. Since all matrices $a \otimes (e_r - e_s)$ have row sum zero, $\mathcal{Z}$ is affinely equivalent to the p-signing zonotope $\mathcal{Z}_A^p$ of A. Therefore, applying the aforementioned enumeration algorithm to $\mathcal{Z}$, we obtain all sign vectors $\sigma = (\sigma_{r,s}^a) \in \{-1,1\}^N$ which correspond to vertices of $\mathcal{Z}$. These are precisely all extremal p-signings of A. We now enumerate all p-partitions $\pi[\sigma]$ associated with extremal p-signings σ in our list. By Theorem 2.4, these are precisely all extremal p-partitions of A. The construction of the sign vectors of $\mathcal{Z}$ dominates the computational complexity of this algorithm. The corollary follows. $\square$

The next corollary concerns the partition problem we started with. The convex functional c on the space of $d \times p$ matrices can be presented by an evaluation oracle that queried on a matrix A^π returns the value $c(A^\pi)$: an oracle algorithm is *strongly polynomial time* if it uses a number of arithmetic operations and oracle queries polynomial in n, and for rational data, runs in time polynomial in the size of the input together with the (rational) oracle answers.

Corollary 3.3 *For every fixed d and $p \geq 2$, the p-partition problem on any given set of n points in d-space and any oracle presented convex functional on $d \times p$-space can be solved in strongly polynomial time using $O(n^{d(p-1)-1})$ arithmetic operations and oracle queries.*

Proof. Let A be a given set of n points in d-space. Suppose first $0 \notin A$. Use the algorithm of Corollary 3.2 to enumerate all extremal p-partitions of A in strongly polynomial time using the claimed number of arithmetic operations. For each extremal p-partition π in the list, form the matrix A^π and query the oracle about the value $c(A^\pi)$. Any extremal p-partition with maximal such value is an optimal solution to the partition problem: indeed, since c is convex, there is always an optimal partition which is extremal. If $0 \in A$ then use the procedure above to obtain an optimal partition of $A \backslash \{0\}$ and place the origin in any part of that partition. The number of queries to the oracle equals the number of extremal p-partitions of A which is bounded in Corollary 3.1, and the computational complexity of all other work is dominated by the enumeration of extremal p-partitions of A which is bounded in Corollary 3.2. □

Finally, we elaborate on the relation between the partition polytopes $\mathcal{P}_A^p$ and the signing zonotope $\mathcal{Z}_A^p$. Recall from Grünbaum (1967) that the *normal cone* $\mathcal{N}(P, F)$ of a convex polytope P in $\mathbb{R}^k$ at its nonempty face F is the relatively open cone of all $c \in \mathbb{R}^k$ for which F is the set of points of P at which the linear functional $\langle c, \cdot \rangle$ is maximized. Note that a nonempty face F of P contains another F' if and only the closure of $\mathcal{N}(P, F')$ contains the closure of $\mathcal{N}(P, F)$. The *normal fan* $\mathcal{N}(P)$ of P is the complex of normal cones of all nonempty faces of P. Note that the entire space $\mathbb{R}^k$ is decomposed into the disjoint union of the cones in $\mathcal{N}(P)$. The normal fan $\mathcal{N}(Z)$ of a polytope Z is a *refinement* of the normal fan $\mathcal{N}(P)$ of a polytope P if the closure of any cone in $\mathcal{N}(P)$ is the union of closures of cones in $\mathcal{N}(Z)$.

Corollary 3.4 *The normal fan $\mathcal{N}(\mathcal{Z}_A^p)$ of the p-signing zonotope $\mathcal{Z}_A^p$ of any set A of nonzero points in $\mathbb{R}^d$ is a refinement of the normal fan $\mathcal{N}(\mathcal{P}_A^p)$ of the p-partition polytope $\mathcal{P}_A^p$ of A. Thus $\mathcal{Z}_A^p$ is the Minkowski sum $\mathcal{Z}_A^p = \lambda \mathcal{P}_A^p + Q$ of a positive multiple of $\mathcal{P}_A^p$ and some polytope Q.*

Proof. The normal cone $\mathcal{N}(\mathcal{P}_A^p, A^\pi)$ of $\mathcal{P}_A^p$ at any vertex A^π is the cone $\mathcal{C}^\pi$ of the extremal partition π, and the normal cone $\mathcal{N}(\mathcal{Z}_A^p, A_\sigma)$ of $\mathcal{Z}_A^p$ at any vertex A_σ is the cone $\mathcal{C}_\sigma$ of the extremal signing σ. Therefore, by Theorem 2.4, the closure of the normal cone of any vertex of $\mathcal{P}_A^p$ is the union of closures of normal cones of vertices $\mathcal{Z}_A^p$. Since, for any polytope, the closure of the normal cone of any face is the intersection of the closures of the normal cones of all vertices contained in that face, it follows that $\mathcal{N}(\mathcal{Z}_A^p)$ is indeed a refinement of $\mathcal{N}(\mathcal{P}_A^p)$.

Now, it is known (cf. Grünbaum (1967)) that the normal fan $\mathcal{N}(Z)$ of a polytope Z is a refinement of the normal fan $\mathcal{N}(P)$ of a polytope P if and only if $Z = \lambda P + Q$ for some positive scalar λ and some polytope Q. This yields the last statement of the corollary. □

References

Alon N., Onn S. 1999. Separable partitions. *Disc. App. Math.* **91** 39–51.

Aviran S., Lev-Tov N., Onn S., Rothblum U.G. (to appear). Vertex characterization of partition polytopes of bipartitions and of planar point sets. *Disc. App. Math.*.

Barnes E.R., Hoffman A.J., Rothblum U.G. 1992. Optimal partitions having disjoint convex and conic hulls. *Math. Prog.* **54** 69–86.

Granot D., Rothblum U.G. 1991. The Pareto set of the partition bargaining game. *Games Econ. Behav.* **3** 163–182.

Gritzmann P., Sturmfels B. 1993. Minkowski addition of polytopes: computational complexity and applications to Gröbner bases. *SIAM J. Disc. Math.* **6** 246–269.

Grünbaum B. 1967. *Convex Polytopes*, Wiley-Interscience, London.

Hwang F.K., Onn S., Rothblum U.G. 2000. Linear shaped partition problems. *Oper. Res. Let.* **26** 159–163.

Hwang F.K., Onn S., Rothblum U.G. 1999. A polynomial time algorithm for shaped partition problems. *SIAM J. Optim.* **10** 70–81.

Onn S., Sturmfels B. 1999. Cutting corners. *Adv. App. Math.* **23** 29–48.

Shmuel Onn
Operations Research, Davidson Faculty of IE & M
Technion - Israel Institute of Technology, 32000 Haifa, Israel.
http://ie.technion.ac.il/~onn
email: onn@ie.technion.ac.il

Leonard J. Schulman
California Institute of Technology, Pasadena California 91125, USA.
email: schulman@caltech.edu