SIAM J. OPTIM.
Vol. 10, No. 1, pp. 70–81

A POLYNOMIAL TIME ALGORITHM FOR SHAPED PARTITION PROBLEMS*

FRANK K. HWANG†, SHMUEL ONN‡, AND URIEL G. ROTHBLUM‡

Abstract. We consider the class of *shaped partition problems* of partitioning n given vectors in d-dimensional criteria space into p parts so as to maximize an arbitrary objective function which is convex on the sum of vectors in each part, subject to arbitrary constraints on the number of elements in each part. This class has broad expressive power and captures NP-hard problems even if either d or p is fixed. In contrast, we show that when both d and p are fixed, the problem can be solved in strongly polynomial time. Our solution method relies on studying the corresponding class of *shaped partition polytopes*. Such polytopes may have exponentially many vertices and facets even when one of d or p is fixed; however, we show that when both d and p are fixed, the number of vertices of any shaped partition polytope is $O(n^{d\binom{p}{2}})$ and all vertices can be produced in strongly polynomial time.

Key words. partition, cluster, optimization, convex, polytope, enumeration, polynomial time, separation, programming

AMS subject classifications. 05A, 15A, 51M, 52A, 52B, 52C, 68Q, 68R, 68U, 90B, 90C

PII. S1052623497344002

1. Introduction.

The *partition problem* concerns the partitioning of vectors $A^1, \ldots, A^n$ in d-space into p parts so as to maximize an objective function which is convex on the sum of vectors in each part; see [3]. Each vector A^i represents d numerical attributes associated with the ith element of the set $[n] = \{1, \ldots, n\}$ to be partitioned. Each ordered partition $\pi = (\pi_1, \ldots, \pi_p)$ of $[n]$ is then associated with the $d \times p$ matrix $A^\pi = \left[\sum_{i\in\pi_1} A^i, \ldots, \sum_{i\in\pi_p} A^i\right]$ whose jth column represents the total attribute vector of the jth part. The problem is to find an *admissible* partition π which maximizes an objective function f given by $f(\pi) = C(A^\pi)$, where C is a real convex functional on $\mathbb{R}^{d\times p}$. Of particular interest is the *shaped partition problem*, where the admissible partitions are those π whose *shape* $(|\pi_1|, \ldots, |\pi_p|)$ lies in a prescribed set Λ of admissible shapes. In this article we concentrate on this later situation.

The shaped partition problem has applications in diverse fields that include circuit layout, clustering, inventory, splitting, ranking, scheduling, and reliability; see [5, 9, 14, 15] and references therein. Further, as we demonstrate later, the problem has expressive power that captures NP-hard problems such as the max-cut problem and the traveling salesman problem, even when the number p of parts or attribute dimension d is fixed.

Our first goal in this article is to demonstrate constructively that a polynomial time algorithm for the shaped partition problem *does* exist when both p and d are fixed. This result is valid when the set Λ of admissible shapes and the function C are

*Received by the editors December 18, 1997; accepted for publication (in revised form) September 17, 1998; published electronically October 20, 1999. The research of the second author was supported in part by the Mathematical Sciences Research Institute at Berkeley, CA, through NSF grant DMS-9022140, by the N. Haar and R. Zinn Research Fund at the Technion, and by the Fund for the Promotion of Research at the Technion. The research of the third author was supported in part by the E. and J. Bishop Research Fund at the Technion and by ONR grant N00014-92-J1142.

http://www.siam.org/journals/siopt/10-1/34400.html

†Department of Applied Mathematics, Chiaotung University, Hsinchu, 30045, Taiwan (fhwang@math.nctu.edu.tw).

‡William Davidson Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, 32000 Haifa, Israel (onn@ie.technion.ac.il, rothblum@ie.technion.ac.il).

presented by oracles. Our first result (formally stated and proved in section 4) is the following:

- *Theorem* 4.2: Any shaped partition problem is solvable in polynomial oracle time using $O(n^{dp^2})$ arithmetic operations and queries.

Our solution method is based on the observation that since C is convex, the shaped partition problem can be embedded into the problem of maximizing C over the *shaped partition polytope* $\mathcal{P}_A^\Lambda$ defined to be the convex hull of all matrices A^π corresponding to partitions of admissible shapes. The class of shaped partition polytopes is very broad and generalizes and unifies classical permutation polytopes such as Birkhoff's polytope and the permutohedron (see, e.g., [8, 19, 21]). Its subclass of *bounded* shaped partition polytopes with lower and upper bounds on the shapes was previously studied in [3], under the assumption that the vectors $A^1, \ldots, A^n$ are *distinct*. Therein a polynomial procedure for testing whether a given A^π is a vertex of $\mathcal{P}_A^\Lambda$ was obtained. This procedure was simplified and extended in [11]. A related but different generalization of classical permutation polytopes, arising when algebraic (representation-theoretical) constraints, rather than shape constraints, are imposed on the permutations involved, was studied in [19] and references therein.

Since a shaped partition polytope is defined as the convex hull of an implicitly presented set whose size is typically exponential in the input size even when both p and d are fixed, an efficient representation as the convex hull of vertices or as the intersection of half-spaces is not readily expected. Our second objective is to prove that, nevertheless, for fixed p and d, the number of vertices of shaped partition polytopes *is* polynomially bounded in n, and that it is possible to explicitly enumerate all vertices in polynomial time. Thus, our second result (formally stated and proved in section 4) is the following:

- *Theorem* 4.3: Any shaped partition polytope $\mathcal{P}_A^\Lambda$ has $O(n^{d\binom{p}{2}})$ vertices which can be produced in polynomial oracle time using $O(n^{d^2p^3})$ arithmetic operations and queries.

An immediate corollary of Theorem 4.3 is that, for fixed d, p, the number of facets of $\mathcal{P}_A^\Lambda$ is polynomially bounded as well and that all facets can be produced in polynomial oracle time (Corollary 4.4). Theorem 4.3 shows, in particular, that it is possible to compute the *number* of vertices efficiently. This might be extendable to the situation of variable d and p, where counting vertices is generally a hard task (cf. [16]), as is counting partitions with various prescribed properties (see [4, 10]). The vertex counting problem for variable d and p will be addressed elsewhere.

A special role in our development is played by *separable* partitions, defined as partitions where vectors in distinct sets are (weakly) separable by hyperplanes. In the special case $d = p = 2$, such partitions had been studied quite extensively (see, e.g., [2, 5, 7, 17]). The case $d = 3, p = 2$ has also been considered quite recently in [6]. Here we study such partitions for all d, p, as well as a class of *generic* partitions, and provide an upper bound on their number and an algorithm for producing them. In our recent related work [1], the precise extremal asymptotical behavior of such partitions is determined.

The embedding of the partition problem into the problem of maximizing the convex function C over the partition polytope is useful due to the optimality of vertices in the latter problem. When Λ consists of a single shape, the optimality of vertices holds for the more general class of asymmetric Schur convex functions, introduced in [13]; see [8]. All of our results apply with C as any asymmetric Schur convex function and Λ consisting of a single shape.

The article is organized as follows. In the next section we formally define the shaped partition problem and shaped partition polytope. We demonstrate the expressive power of this problem by giving four examples. For the first two examples, in which the parameters d, p are typically small and fixed, Theorem 4.2 provides a polynomial time solution. The last two examples show that the max-cut problem and traveling salesman problem can be modeled as shaped partition problems with fixed $p = 2$ and $d = 1$, respectively, and that the corresponding polytopes have exponentially many vertices. In section 3 we study separability properties of vertices of shaped partition polytopes and discuss separable and generic partitions. In the final section, section 4, we use our preparatory results of section 3 to establish Theorems 4.2 and 4.3 and Corollary 4.4.

2. Shaped partition problems and polytopes. A *p-partition* of $[n] := \{1, \ldots, n\}$ is an ordered collection $\pi = (\pi_1, \ldots, \pi_p)$ of p disjoint sets (possibly empty) whose union is $[n]$. A *p-shape* of n is a tuple $\lambda = (\lambda_1, \ldots, \lambda_p)$ of nonnegative integers $\lambda_1, \ldots, \lambda_p$ satisfying $\sum_{i=1}^{p} \lambda_i = n$. The *shape of a p-partition* π is the p-shape of n given by $|\pi| := (|\pi_1|, \ldots, |\pi_p|)$. If Λ is a set of p-shapes of n, then a Λ-*partition* is any partition π whose shape $|\pi|$ is a member of Λ.

Let A be a real $d \times n$ matrix; for $i = 1, \ldots, n$, we use A^i to denote the ith column of A. For each p-partition π of $[n]$ we define the *A-matrix* of π to be the $d \times p$ matrix

$$A^\pi = \left[\sum_{i \in \pi_1} A^i, \ldots, \sum_{i \in \pi_p} A^i \right],$$

with $\sum_{i \in \pi_j} A^i := 0$ when $\pi_j = \emptyset$. We consider the following algorithmic problem.

Shaped Partition Problem. Given positive integers d, p, n, matrix $A \in \mathbb{R}^{d \times n}$, the nonempty set Λ of p-shapes of n, and the objective function on Λ-partitions given by $f(\pi) = C(A^\pi)$ with C convex on $\mathbb{R}^{d \times p}$, find a Λ-partition π^* that maximizes f and, specifically, satisfies

$$f(\pi^*) = \max\{f(\pi) : \ |\pi| \in \Lambda\}.$$

Of course, the complexity of this problem depends on the presentation of Λ and C, but we will construct algorithms that work in strongly polynomial time and can cope with minimal information on Λ and C. Specifically, we assume that the set of admissible p-partitions Λ can be represented by a membership oracle that, on query λ, answers whether $\lambda \in \Lambda$. The convex functional C on $\mathbb{R}^{d \times p}$ can be presented by an evaluation oracle that, on query A^π with π a Λ-partition, returns $C(A^\pi)$.

Since C is convex, the shaped partition problem can be embedded into the problem of maximizing C over the convex hull of A-matrices of feasible partitions, defined as follows.

Shaped Partition Polytope. For a matrix $A \in \mathbb{R}^{d \times n}$ and nonempty set Λ of p-shapes of n, we define the *shaped partition polytope* $\mathcal{P}_A^\Lambda$ to be the convex hull of all A-matrices of Λ-partitions, that is,

$$\mathcal{P}_A^\Lambda := \operatorname{conv}\{A^\pi : \ |\pi| \in \Lambda\} \subset \mathbb{R}^{d \times p}.$$

We point out that for any A, the polytope $\mathcal{P}_A^\Lambda$ is the image of the shaped partition polytope P_I^Λ, with I the $n \times n$ identity, under the projection $X \mapsto AX$. In [12] this is exploited, for the situation where the function C is *linear* and $\Lambda = \{\lambda : \ l \leq \lambda \leq u\}$ is a set of *bounded* shapes, to solve the corresponding shaped partition problem for all n, d, p in polynomial time by linear programming over P_I^Λ.

We now demonstrate the expressive power of the shaped partition problem. In particular, we show that even if one of d or p is fixed, the shaped partition problem may be NP-hard, and the number of vertices of the shaped partition polytope may be exponential. Therefore, polynomial time algorithms for optimization and vertex enumeration are expected to (and, as we show, *do*) exist only when *both* d and p are fixed. We start with two examples in which it is natural to have d and p small and fixed.

EXAMPLE 2.1 (splitting). *The n assets of a company are to be split among its p owners as follows. For $j = 1, \ldots, p$, the jth owner prescribes a nonnegative vector $A_j = (A_{j,1}, \ldots, A_{j,n})$ with $\sum_{j=1}^n A_{i,j} = 1$, whose entries represent the relative values of the various assets to this owner. A partition $\pi = (\pi_1, \ldots, \pi_p)$ is sought which splits the assets among the owners and maximizes the l_q-norm $(\sum_{j=1}^p |\sum_{i\in\pi_j} A_{j,i}|^q)^{\frac{1}{q}}$ of the total value vector whose jth entry $\sum_{i\in\pi_j} A_{j,i}$ is the total relative value of the assets allocated to the jth owner by π. An alternative interpretation of the splitting problem concerns the division of an estate consisting of n assets among p inheritors having equal rights against the estate. With $p = 2$, the model captures a problem of a divorcing couple dividing their joint property* [5, 9].

FORMULATION: $n,\ d = p,\ A = (A_{j,i}),\ \Lambda = \{\textit{All } p\textit{-shapes}\},\ f(\pi) = C(A^\pi)$ *with*

$$C : \mathbb{R}^{p\times p} \longrightarrow \mathbb{R} :\ M \mapsto \sum_{i=1}^{p} |M_{i,i}|^q.$$

For fixed p, Theorem 4.2 *asserts that we can find an optimal partition in polynomial time $O(n^{p^3})$, while the number p^n of Λ-partitions is exponential. We note that other (convex) functions C can be used within our framework. In particular, if C is linear on $\mathbb{R}_+^{p\times p}$, e.g., when $q = 1$, our results of* [12] *apply and yield a polynomial time solution even when p is variable.*

EXAMPLE 2.2 (balanced clustering). *Given are $n = pm$ objects represented by attribute vectors $A^1, \ldots, A^n \in \mathbb{R}^d$. The objects are to be grouped in p clusters, each containing m points, so as to minimize the sum of cluster variance of a partition π given by $\sum_{i=1}^p \left(\frac{1}{|\pi_i|}\sum_{j\in\pi_i} ||A^j - \bar{A}^{\pi_i}||^2\right)$, where $||\cdot||$ denotes the l_2-norm and $\bar{A}^{\pi_i} := \frac{1}{|\pi_i|}\sum_{j\in\pi_i} A^j$ is the barycenter of the ith cluster.*

FORMULATION: $n = pm,\ d,\ p,\ A = (A^1, \ldots, A^n),\ \Lambda = \{m^p = (m, \ldots, m)\},\ f(\pi) = C(A^\pi)$ *with*

$$C : \mathbb{R}^{d\times p} \longrightarrow \mathbb{R} :\ M \mapsto ||M||^2 = \sum_{i=1}^{d}\sum_{j=1}^{p} M_{i,j}^2.$$

Here, we use the fact that $f(\pi) = \frac{1}{m^2}\sum_{i=1}^n ||A^i||^2 - \frac{1}{m^2}\sum_{j=1}^p ||\sum_{i\in\pi_j} A^i||^2$. For fixed d, p, by Theorem 4.2 *we can find an optimal balanced clustering in polynomial time $O(n^{dp^2})$, while the number of Λ-partitions is exponential $\Omega(p^n n^{\frac{1-p}{2}})$.*

The next two examples show that unless *both* d and p are fixed, the shaped partition problem may be NP-hard. The idea is simple: the formulation is such that every Λ-partition π gives a distinct vertex A^π of the shaped partition polytope $\mathcal{P}_A^\Lambda$. Then, *any* function f on Λ-partitions factors as $f(\pi) := C(A^\pi)$ for suitable convex C

on $\mathcal{P}_A^\Lambda$, say, the one given by

$$C(X) := \inf\left\{\sum_{|\pi|\in\Lambda}\theta_\pi f(\pi):\ \sum_\pi \theta_\pi A^\pi = X,\ \sum_\pi \theta_\pi = 1,\ \theta_\pi \geq 0\right\}.$$

In the following examples, the membership oracle for Λ and the evaluation oracle for $f(\pi) := C(A^\pi)$, restricted to A-matrices, are easily polynomial time realizable from the natural data for the problem.

EXAMPLE 2.3 (max-cut problem and unit cube). *Find a cut with maximum number of crossing edges in a given graph $G = ([n], E)$.*

FORMULATION: $n = d,\ p = 2,\ A = I_n,\ \Lambda = \{$*all* 2-*shapes*$\}$,

$$f(\pi) = \#\{e \in E:\ |e \cap \pi_1| = 1\}.$$

Here, the A-matrices of Λ-partitions are precisely all $(0,1)$-valued $n \times 2$ matrices with each row sum equal to 1; in particular, each such matrix is determined by its first column. It follows that the shaped partition polytope $\mathcal{P}_A^\Lambda$ has 2^n vertices that stand in bijection with Λ-partitions and is affinely equivalent to the n-dimensional unit cube by projection of matrices onto their first column. So, each A^π is a distinct vertex of $\mathcal{P}_A^\Lambda$ and there is a convex C on $\mathbb{R}^{d\times 2}$ such that $f(\pi) = C(A^\pi)$ for all π.

EXAMPLE 2.4 (traveling salesman problem and permutohedron). *Find a shortest Hamiltonian path on n sites under a given symmetric nonnegative matrix D, where $D_{i,j}$ represents the distance between sites i and j.*

FORMULATION: $n = p,\ d = 1,\ A = (1, \ldots, n),\ \Lambda = \{1^n = (1, \ldots, 1)\}$,

$$f(\pi) = -\sum_{j=1}^{n-1} D_{\pi_j,\pi_{j+1}},$$

where we regard a partition simply as the corresponding permutation. The matrices A^π in this case are simply all permutations of A. The shaped partition polytope $\mathcal{P}_A^\Lambda$ has $n!$ vertices that stand in bijection with Λ-partitions, and is the so-called permutohedron. Since each A^π is a distinct vertex of $\mathcal{P}_A^\Lambda$, there is again a convex C on $\mathbb{R}^n$ such that $f(\pi) = C(A^\pi)$ for all π.

3. Vertices and generic partitions.

In this section we show that every vertex of any shaped partition polytope $\mathcal{P}_A^\Lambda$ equals the A-matrix A^π of some A-*generic* partition, a notion that we introduce and develop below.

The convex hull of a subset U in $\mathbb{R}^d$ will be denoted $\mathrm{conv}(U)$. Two finite sets U, V of points in $\mathbb{R}^d$ are *separable* if there is a vector $h \in \mathbb{R}^d$ such that $h^T u < h^T v$ for all $u \in U$ and $v \in V$ with $u \neq v$; in this case, we refer to h as a *separating vector* of U and V. The proof of the following characterization of separability is standard and is left to the reader. It implies, in particular, that if U and V are separable, then $|U \cap V| \leq 1$.

LEMMA 3.1. *Let U and V be finite sets of $\mathbb{R}^d$. Then U and V are separable if and only if their convex hulls either are disjoint or intersect in a single point that is a common vertex of both.*

Let A be a given $d \times n$ matrix. For a subset $S \subseteq [n]$, let $A^S = \{A^i:\ i \in S\}$ be the set of columns of A indexed by S (with multiple copies of columns identified). A p-partition $\pi = (\pi_1, \ldots, \pi_p)$ is A-*separable* if the sets A^{π_r} and A^{π_s} are separable for each pair $1 \leq r < s \leq p$, that is, if for each pair $1 \leq r < s \leq p$ there is a vector

$h_{r,s} \in \mathbb{R}^d$ such that $h_{r,s}^T A^i < h_{r,s}^T A^j$ for all $i \in \pi_r$ and $j \in \pi_s$ with $A^i \neq A^j$. We have the following lemma, which generalizes a result of [3] from matrices with no zero columns and no repeated columns.

LEMMA 3.2. *Let A be a matrix in $\mathbb{R}^{d\times n}$, let Λ be a nonempty set of p-shapes of n, and let π be a Λ-partition. If A^π is a vertex of $\mathcal{P}_A^\Lambda$, then π is an A-separable partition.*

Proof. The claim being obvious for $p = 1$, suppose that $p \geq 2$. Let A^π be a vertex of $\mathcal{P}_A^\Lambda$. Then there is a matrix $C \in \mathbb{R}^{d\times p}$ such that the linear functional on $\mathbb{R}^{d\times p}$ given by the inner product $\langle C, X\rangle = \sum_{i=1}^d \sum_{j=1}^p C_i^j X_i^j$ is uniquely maximized over $\mathcal{P}_A^\Lambda$ at A^π. Pick any pair $1 \leq r < s \leq p$, and let $h_{r,s} = C^s - C^r$. Suppose there are $i \in \pi_r$ and $j \in \pi_s$ with $A^i \neq A^j$ (otherwise A^{π_r} and A^{π_s} are trivially separable). Let π' be the Λ-partition obtained from π by switching i and j, i.e., taking $\pi'_r := \pi_r \cup \{j\} \setminus \{i\}$, $\pi'_s := \pi_s \cup \{i\} \setminus \{j\}$, and $\pi'_t := \pi_t$ for all $t \neq r, s$. Then $(A^{\pi'})^r = (A^\pi)^r + A^j - A^i \neq (A^\pi)^r$, and hence $A^{\pi'} \neq A^\pi$. By the choice of C, we have $\langle C, A^{\pi'}\rangle < \langle C, A^\pi\rangle$ and so

$$h_{r,s}^T(A^j - A^i) = (C^s - C^r)^T(A^j - A^i) = \sum_{t=1}^p (C^t)^T((A^\pi)^t - (A^{\pi'})^t) = \langle C, A^\pi - A^{\pi'}\rangle > 0.$$

This proves that A^{π_r}, A^{π_s} are separable for each pair $1 \leq r < s \leq p$, hence π is A-separable. ∎

We need some more terminology. Let $A \in \mathbb{R}^{d\times n}$. A p-partition $\pi = (\pi_1, \ldots, \pi_p)$ of $[n]$ is A-*disjoint* if $\mathrm{conv}(A^{\pi_r})$ and $\mathrm{conv}(A^{\pi_s})$ are disjoint for each pair $1 \leq r < s \leq p$. As the convex hulls of finite sets are disjoint if and only if the sets can be strictly separated by a hyperplane, we have that π is A-disjoint if and only if for each pair $1 \leq r < s \leq n$ there exists a vector $h_{r,s} \in \mathbb{R}^d$ such that $(h_{r,s})^T A^i < (h_{r,s})^T A^S$ for all $i \in \pi_r$ and $j \in \pi_s$. Of course, A-disjointness implies A-separability, and the two properties coincide when the columns of A are distinct.

For $v \in \mathbb{R}^d$ denote by $\bar{v} \in \mathbb{R}^{d+1}$ the vector obtained by appending a first coordinate 1 to v. For a matrix $A \in \mathbb{R}^{d\times n}$ and indices $1 \leq i_0 < \cdots < i_d \leq n$, denote

$$\mathrm{sign}_A(i_0, \ldots, i_d) := \mathrm{sign}(\det[\bar{A}^{i_0}, \ldots, \bar{A}^{i_d}]) \in \{-1, 0, 1\}.$$

A matrix A is *generic* if its columns are in affine general position, that is, if any set of $d + 1$ vectors or less from among $\{\bar{A}^i : i = 1, \ldots, n\}$ are linearly independent; in particular, if $n > d$ this is the case if and only if all signs $\mathrm{sign}_A(i_0, \ldots, i_d)$ for indices $1 \leq i_0 < i_1 < \cdots < i_d \leq n$ are nonzero. Also, the columns of a generic matrix are distinct.

We next provide a representation of the set of A-disjoint 2-partitions for generic matrices A. The case where $n \leq d$ is simple.

LEMMA 3.3. *Let $A \in \mathbb{R}^{d\times n}$ be generic, $p \leq 2$, and $n \leq d$. Then every p-partition of $[n]$ is A-disjoint.*

Proof. It suffices to consider the case $p = 2$. A standard result from linear algebra shows that as $\bar{A}^1, \ldots, \bar{A}^n$ are linearly independent, the range of $[\bar{A}^1, \ldots, \bar{A}^n]^T$ is $\mathbb{R}^n$. Hence, given a 2-partition π of $[n]$, there is a vector $\mu \in \mathbb{R}^{d+1}$ with $\mu^T A^i > 0$ for each $i \in \pi$ and $\mu^T A^j < 0$ for each $j \in \pi_2$; with C obtained from μ by truncating its first coordinate μ_1, we then have $C^T A^i > -\mu_1 > C^T A^j$ for all $i \in \pi_1$ and $j \in \pi_2$, proving that π is A-disjoint. ∎

Let $A \in \mathbb{R}^{d\times n}$ be generic with $n \geq d$. For any d-subset $I = \{i_1, \ldots, i_d\}$ of $[n]$ with $i_1 < \cdots < i_d$, define

$$I_A^- := \{i_0 \in [n] : \mathrm{sign}_A(i_0, i_1, \ldots, i_d) = -1\}, \quad I_A^+ := \{i_0 \in [n] : \mathrm{sign}_A(i_0, i_1, \ldots, i_d) = 1\}.$$

Of course, $\{I_A^-, I_A^+\}$ is a 2-partition of $[n]\backslash I$. Let $I \subseteq [n]$ be a d-set and (J^-, J^+) be a 2-partition of I. The 2-partitions of $[n]$ *associated* with A, I, and (J^-, J^+) are defined to be either of the two 2-partitions $\pi^- := (I_A^- \cup J^-, I_A^+ \cup J^+)$ and $\pi^+ := (I_A^+ \cup J^+, I_A^- \cup J^-)$.

LEMMA 3.4. *Let $A \in \mathbb{R}^{d \times n}$ be generic, with $n \geq d$. Then the set of A-disjoint 2-partitions is the set of all 2-partitions associated with A, d-sets $I \subseteq [n]$ and 2-partitions (J^-, J^+) of I.*

Proof. We will show that for each d-set $I \subseteq [n]$ and 2-partition (J^-, J^+) of I, the two 2-partitions associated with A, I, and (J^-, J^+) are A-disjoint and that each A-disjoint 2-partition is generated in this way.

First, let $I \subseteq [n]$ have d-elements, say, $i_1 < \cdots < i_d$, and let (J^-, J^+) be a 2-partition of I. Then $H := \{x \in \mathbb{R}^d : \det[\bar{x}, \bar{A}^{i_1}, \ldots, \bar{A}^{i_d}] = 0\}$ is a hyperplane that contains the columns of A indexed by I; this hyperplane can be written as $\{x \in \mathbb{R}^d : h^T x = \gamma\}$ for some $h \in \mathbb{R}^d$ and $\gamma \in \mathbb{R}$ such that $I_A^- = \{i \in [n] : h^T A^i < \gamma\}$ and $I_A^+ = \{i \in [n] : h^T A^i > \gamma\}$. Thus, $h^T A^i < h^T A^U < h^T A^j$ for all $i \in I_A^-, u \in I$, and $j \in I_A^+$. We next observe that $B = [A^{i_1}, \ldots, A^{i_d}]$ is generic, hence Lemma 3.3 ensures that the 2-partition $\{j : i_j \in J_-\}, \{j : i_j \in J_+\}$ of $[d]$ is B-disjoint. Thus, there exists a vector $d \in \mathbb{R}^d$ with $d^T A^i > d^T A^j$ for all $i \in J^-$ and $j \in J^+$. For sufficiently small positive t, we then have that $(C + td)A^i < (C + td)^T A^j$ for all $i \in I_A^- \cup J^-$ and $j \in I_A^+ \cup J^+$, proving that $(I_A^- \cup J^-, J_A^+ \cup J^+)$ is A-disjoint. It follows immediately that $(I_A^+ \cup J^+, I_A^- \cup J^-)$ is A-disjoint too, proving that the two 2-partitions of $[n]$ associated with A, I and the 2-partition (J_-, J_+) of I are A-disjoint.

Next assume that π is an A-disjoint 2-partition. Then there exists a hyperplane strictly separating A^{π_1} and A^{π_2}. Any such hyperplane can be perturbed to a hyperplane that is spanned by d columns of A and weakly separates A^{π_1} and A^{π_2} (the details of constructing such a perturbation are left to the reader). In particular, if $A^{i_1}, \ldots, A^{i_d}$ span the hyperplane and $1 \leq i_1 < \cdots \leq i_d \leq n$, then for $I := \{i_1, \ldots, i_d\}$ either $\pi_1^1 \subseteq I_A^- \cup I$ and $\pi_2 \subseteq I_A^+ \cup I$ or $\pi_1 \subseteq I_A^+ \cup I$ and $\pi_2 \subseteq I_A^- \cup I$. In the former case we have $\pi = (I_A^- \cup J^-, I_A^+ \cup J^+)$ for $J^- = \pi_1 \cap I$ and $J^+ = \pi_2 \cap I$, and in the latter case $\pi = (I_A^+ \cup J^+, I_A^- \cup J^-)$ for $J^+ = \pi_1 \cap I$ and $J^- = \pi_2 \cap J$. ◻

Let $p \geq 2$. With each list $[\pi^{r,s} = (\pi_1^{r,s}, \pi_2^{r,s}) : 1 \leq r < s \leq p]$ of $\binom{p}{2}$ 2-partitions of $[n]$ *associate a* p-tuple $\pi = (\pi_1, \ldots, \pi_p)$ of subsets of $[n]$ as follows: for $r = 1, \ldots, p$ put

$$\pi_r := \left(\cap_{j=r+1}^{p} \pi_1^{r,j}\right) \bigcap \left(\cap_{j=1}^{r-1} \pi_2^{j,r}\right).$$

Since $\pi_r \subseteq \pi_1^{r,s}$ and $\pi_s \subseteq \pi_2^{r,s}$ for all $1 \leq r < s \leq p$, the elements π_i of the p-tuple associated with the given list are pairwise disjoint. If $\cup_{i=1}^{p} \pi_i = [n]$ holds as well, then π is a p-partition that will be called the partition *associated* with the given list.

LEMMA 3.5. *For $A \in \mathbb{R}^{d \times n}$ and $p \geq 2$, the set of A-disjoint p-partitions equals the set of p-partitions associated with lists of $\binom{p}{2}$ A-disjoint 2-partitions.*

Proof. First, consider a p-partition π associated with a list of $\binom{p}{2}$ A-disjoint 2-partitions. Then, for each $1 \leq r < s \leq p$,

$$\operatorname{conv}(A^i : i \in \pi_r) \cap \operatorname{conv}(A^i : i \in \pi_s) \subseteq \operatorname{conv}(A^i : i \in \pi_1^{r,s}) \cap \operatorname{conv}(A^i : i \in \pi_2^{r,s}) = \emptyset,$$

so π is A-disjoint. Conversely, let $\pi = (\pi_1, \ldots, \pi_p)$ be an A-disjoint p-partition. Consider any pair $1 \leq r < s \leq p$. Since $\operatorname{conv}(A^{\pi_r})$ and $\operatorname{conv}(A^{\pi_s})$ are disjoint, there is a hyperplane $H_{r,s}$ that contains no column of A and defines two corresponding half-spaces $H_{r,s}^-$ and $H_{r,s}^+$ that satisfy $A^{\pi_r} \subset H_{r,s}^-$ and $A^{\pi_s} \subset H_{r,s}^+$. Let $\pi^{r,s} := (\pi_1^{r,s}, \pi_2^{r,s})$ be the A-disjoint 2-partition defined by $\pi_1^{r,s} := \{i \in [n] : A^i \in H_{r,s}^-\}$ and $\pi_2^{r,s} := \{i \in$

$[n] : A^i \in H^+_{r,s}\}$. Let π' be the p-tuple associated with the constructed $\pi^{r,s}$'s. Then the sets of π' are pairwise disjoint, and for $i = 1, \ldots, p$, we have

$$\pi_i \subseteq \left(\cap_{j=i+1}^{p} \pi_1^{i,j}\right) \bigcap \left(\cap_{j=1}^{i-1} \pi_2^{j,i}\right) = \pi_i'.$$

Since $[n] = \cup_{i=1}^p \pi_i \subseteq \cup_{i=1}^p \pi_i'$, it follows that $\pi = \pi'$ is the p-partition associated with the constructed list of $\binom{p}{2}$ A-disjoint 2-partitions. ∎

For each $\epsilon > 0$ define the *ϵ-perturbation* $A(\epsilon) \in \mathbb{R}^{d\times n}$ *of* A as follows: for $i = 1, \ldots, n$, let the ith column of $A(\epsilon)$ be $A(\epsilon)^i := A^i + \epsilon M_d^i$, where $M_d^i := [i, i^2, \ldots, i^d]^T$ is the image of i on the moment curve in $\mathbb{R}^d$. Consider any $1 \le i_0 < \cdots < i_d \le n$. Then the determinant

$$D(\epsilon) := \det[\bar{A}(\epsilon)^{i_0}, \ldots, \bar{A}(\epsilon)^{i_d}] = \sum_{j=0}^{d} D_j \epsilon^j$$

is a polynomial of degree d in ϵ, with D_d being the Van der Monde determinant $\det[\bar{M}_d^{i_0}, \ldots, \bar{M}_d^{i_d}]$, which is known to be nonzero. So for all sufficiently small $\epsilon > 0$, $\text{sign}_{A(\epsilon)}(i_0, \ldots, i_d) = \text{sign}(D(\epsilon))$ equals the sign of the first nonzero coefficient among $D_0, \ldots, D_d$ and is either -1 or 1 and independent of ϵ. We define the *generic sign* of A at $(i_0, \ldots, i_d)$, denoted $\chi_A(i_0, \ldots, i_d)$, as the common value of $\text{sign}_{A(\epsilon)}(i_0, \ldots, i_d)$ for all sufficiently small positive ϵ.

Lemma 3.6. *Let $A \in \mathbb{R}^{d\times n}$ and $p \ge 1$. For all sufficiently small $\epsilon > 0$, $A(\epsilon)$ is generic and the set of $A(\epsilon)$-disjoint p-partitions is the same. Further, for every d-set $I \in [n]$, the sets $I^-_{A(\epsilon)}$ and $I^+_{A(\epsilon)}$ are independent of ϵ.*

Proof. By Lemma 3.5, the set of $A(\epsilon)$-disjoint p-partitions is entirely determined by the set of $A(\epsilon)$-disjoint 2-partitions. Thus, it suffices to consider only $p = 2$.

First assume that $n < d$. In this case augment A with $n + 1 - d$ zero vectors to obtain a matrix $A' \in \mathbb{R}^{d\times(d+1)}$. The above arguments show that for sufficiently small positive ϵ, $\det \bar{A}'(\epsilon)$ is nonzero, implying that $\bar{A}(\epsilon)^1, \ldots, \bar{A}(\epsilon)^n$ are linearly independent. From Lemma 3.3 it follows that for such ϵ, the set of $A(\epsilon)$-disjoint 2-partitions of $[n]$ is the set of all 2-partitions of $[n]$.

Next assume that $n > d$. As explained above, for all sufficiently small $\epsilon > 0$, $\text{sign}_{A(\epsilon)}\ (i_0, \ldots, i_d)$ equals the nonzero generic sign $\chi_A(i_0, \ldots, i_d)$ for all $1 \le i_0 < \cdots < i_d \le n$. It follows that for all sufficiently small ϵ, the matrix $A(\epsilon)$ is generic, and for every d-set I, the sets $I^-_{A(\epsilon)}$ and $I^+_{A(\epsilon)}$ are independent of ϵ. By Lemma 3.4, the set of $A(\epsilon)$-disjoint 2-partitions is the set of all pairs of 2-partitions of $[n]$ associated with A, d-sets $I \subseteq [n]$, and 2-partitions (J^-, J^+) of I; but each such pair depends only on $I^-_{A(\epsilon)}, I^+_{A(\epsilon)}$, J^-, and J^+. Hence the set of $A(\epsilon)$-disjoint 2-partitions is the same for all sufficiently small $\epsilon > 0$. ∎

Let $A \in \mathbb{R}^{d\times n}$. A p-partition of $[n]$ is *A-generic* if it is $A(\epsilon)$-disjoint for all sufficiently small $\epsilon > 0$. Denote by Π_A^p the set of A-generic p-partitions.

Lemma 3.6 shows that for all sufficiently small $\epsilon > 0$, the set of $A(\epsilon)$-disjoint partitions is the same and equals Π_A^p. The final lemma of this section links vertices of shaped partition polytopes with generic partitions.

Lemma 3.7. *Let $A \in \mathbb{R}^{d\times n}$ and let Λ be a nonempty set of p-shapes of $[n]$. Then every vertex of the polytope $\mathcal{P}_A^\Lambda$ has a representation as the A-matrix A^π of some A-generic Λ-partition.*

Proof. Let $B \in \mathbb{R}^{d\times p}$ be a vertex of $\mathcal{P}_A^\Lambda$ and let $C \in \mathbb{R}^{d\times p}$ be a matrix such that $\langle C, \cdot\rangle$ is uniquely maximized over $\mathcal{P}_A^\Lambda$ at B. Let $\Pi := \{\pi : |\pi| \in \Lambda\}$ be the set of

Λ-partitions and let $\Pi^* := \{\pi \in \Pi : A^\pi = B\}$. Then there is a sufficiently small $\epsilon > 0$ such that $\langle C, A(\epsilon)^{\pi^*}\rangle > \langle C, A(\epsilon)^{\pi}\rangle$ for all $\pi^* \in \Pi^*$ and $\pi \in \Pi \setminus \Pi^*$, and in addition, as guaranteed by Lemma 3.6, $A(\epsilon)$ is generic and the set of $A(\epsilon)$-disjoint p-partitions equals Π^p_A. For such ϵ, $\langle C, \cdot\rangle$ is maximized over the perturbed polytope $P^\Lambda_{A(\epsilon)}$ at a vertex of the form $A(\epsilon)^{\pi^*}$ for some $\pi^* \in \Pi^*$. By Lemma 3.2, π^* is $A(\epsilon)$-separable. Since $A(\epsilon)$ is generic it has distinct columns, and therefore π^* is also $A(\epsilon)$-disjoint. We conclude that π^* is A-generic, proving that π^* contains a generic partition. ∎

4. Optimization and vertex enumeration. We now use the facts established in the previous section to prove our main results. Our computational complexity terminology is fairly standard (cf. [20]). In all our algorithms, the positive integer n will be input in *unary* representation, whereas all other numerical data such as the matrix A will be input in *binary* representation. An algorithm is *strongly polynomial time* if it uses a number of arithmetic operations polynomially bounded in n, and runs in time polynomially bounded in n plus the bit size of all other numerical input.

LEMMA 4.1. *Let d, p be fixed. For any $A \in \mathbb{R}^{d\times n}$, the set Π^p_A of A-generic p-partitions has $|\Pi^p_A| = O(n^{d\binom{p}{2}})$. Further, there is an algorithm that, given $n \in \mathbb{N}$ and $A \in \mathbb{Q}^{d\times n}$, produces Π^p_A in strongly polynomial time using $O(n^{dp^2})$ arithmetic operations.*

Proof. If $n \le d$, the set of A-generic p-partitions is the set of all partitions, of which there are $p^n \le p^d$. Henceforth we assume that $n > d$. If $p = 1$, then $\Pi^p_A := \{([n])\}$ consists of the single p-partition $([n])$. Suppose now that $p \ge 2$. For each choice $1 \le i_0 < \cdots < i_d \le n$, compute the generic sign $\chi_A(i_0, \ldots, i_d)$ as follows. Evaluate the polynomial

$$D(\epsilon) := \det[\bar{A}(\epsilon)^{i_0}, \ldots, \bar{A}(\epsilon)^{i_d}] = \sum_{j=0}^{d} D_j \epsilon^j$$

at $\epsilon = 0, 1, \ldots, d$ to obtain $D(0), D(1), \ldots, D(d)$. Each evaluation involves the computation of the determinant of a matrix of order $d+1$ and can be done, say, by Gaussian elimination, using $O(d^3)$ arithmetic operations and, for rational A, in strongly polynomial time. Then, solve the following linear system of equations:

$$\sum_{j=0}^{d} \epsilon^j D_j = D(\epsilon), \quad \epsilon = 0, \ldots, d,$$

to obtain the indeterminates $D_0, \ldots, D_d$. This can be done by inverting the nonsingular Vandermonde matrix of coefficients of this system, again by Gaussian elimination. The generic sign $\chi_A(i_0, \ldots, i_d)$ is then the sign of the first nonzero D_i. So, for fixed d, the number of arithmetic operations needed to compute all $\binom{n}{d+1}$ generic signs is $O(\binom{n}{d+1} d^4) = O(n^{d+1})$.

By Lemma 3.6, for sufficiently small positive ϵ, for each d-set $I \subseteq [n]$, $I^-_{A(\epsilon)}$ and $I^+_{A(\epsilon)}$ are independent of ϵ. For a d-set $I \subseteq [n]$ and such ϵ, $I^-_{A(\epsilon)}$ and $I^+_{A(\epsilon)}$ are available from the above signs that determine $\det[\bar{A}^i, \bar{A}^{i_1}, \ldots, \bar{A}^{i_q}]$ for each $i \in [n] \setminus J$ (a permutation that puts $\bar{A}^i$ into the right location may be applied). Further, from Lemmas 3.6 and 3.4, Π^2_A equals the common set of $A(\epsilon)$-disjoint partitions for sufficiently small positive ϵ, and this set is the set of partitions of $[n]$ of the form $(I^-_{A(\epsilon)} \cup J^-, I^+_{A(\epsilon)} \cup J^+)$ or $(I^+_{a(\epsilon)} \cup J^+, I^-_{A(\epsilon)} \cup J^-)$, where I is a d-subset of $[n]$ and

(J^-, J^+) is a 2-partition of I. For each d-set $I \subseteq [n]$, the common 2-partitions $(I^-_{A(\epsilon)}, I^+_{A(\epsilon)})$ for sufficiently small positive ϵ have been determined; hence a list of the 2-partitions in Π^2_A is available (the construction may contain duplicates). As there are $\binom{n}{d}$ d-subsets I and 2^d 2-partitions (J^-, J^+) of each I, we have $|\Pi^2_A| \le 2^{d+1}\binom{n}{d} = O(n^d)$ and all partitions in Π^2_A can be obtained from the generic signs, again using $O(n^{d+1})$ operations.

For sufficiently small positive ϵ, Π^p_A is the common set of $A(\epsilon)$-disjoint p-partitions and Π^2_A is the common set of $A(\epsilon)$-disjoint 2-partitions. It follows from Lemma 3.5 that Π^p_A is the set of all p-partitions associated with lists of $\binom{p}{2}$ 2-partitions from Π^2_A. This shows that

$$|\Pi^p_A| \le |\Pi^2_A|^{\binom{p}{2}} = O(n^{d\binom{p}{2}}).$$

To construct Π^p_A, produce all such lists of $\binom{p}{2}$ 2-partitions from Π^2_A; for each list, form the associated p-tuple π and test if it is a partition (i.e., if $\cup_{i=1}^p \pi_i = [n]$). As there are $O(n^{d\binom{p}{2}})$ lists, all this work can be done easily using $O(n^{dp^2})$ arithmetic operations, which subsumes the work for computing the generic signs and constructing Π^2_A, and is the claimed bound. ∎

We can now provide the solution of the shaped partition problem. The set of admissible p-partitions Λ can be represented by a membership oracle that, on query, λ answers whether $\lambda \in \Lambda$. The convex functional C on $\mathbb{R}^{d\times p}$ can be presented by an evaluation oracle that, on query A^π with π a Λ-partition, returns $C(A^\pi)$. The oracle for C will be called *M-guaranteed* if $C(A^\pi)$ is guaranteed to be a rational number whose absolute value is no larger than M for any Λ-partition π. The algorithm is then *strongly polynomial oracle time* if it uses a number of arithmetic operations and oracle queries polynomially bounded in n and runs in time polynomially bounded in n plus the bit size of A and M.

THEOREM 4.2. *For every fixed d, p, there is an algorithm that, given $n, M \in \mathbb{N}$, $A \in \mathbb{Q}^{d\times n}$, oracle-presented nonempty set Λ of p-shapes of n, and M-guaranteed oracle-presented convex functional C on $\mathbb{Q}^{d\times p}$, solves the shaped partition problem in strongly polynomial oracle time using $O(n^{dp^2})$ arithmetic operations and oracle queries.*

Proof. Use the algorithm of Lemma 4.1 to construct the set Π^p_A of A-generic p-partitions in strongly polynomial time using $O(n^{dp^2})$ arithmetic operations. Then test shapes of the partitions in the list to obtain the subset $\Pi^\Lambda := \{\pi \in \Pi^p_A : |\pi| \in \Lambda\}$ of A-generic Λ-partitions by querying the Λ-oracle on each of the $|\Pi^p_A| = O(n^{d\binom{p}{2}})$ partitions in Π^p_A. Since C is convex, it is maximized over the shaped partition polytope $\mathcal{P}^\Lambda_A$ at a vertex of $\mathcal{P}^\Lambda_A$. By Lemma 3.7, this vertex equals the A-matrix A^π of some partition in Π^Λ. Therefore, any $\pi^* \in \Pi^\Lambda$ achieving $C(A^{\pi^*}) = \max\{C(A^\pi) : \pi \in \Pi^\Lambda\}$ is an optimal solution to the shaped partition problem. To find such π^*, compute for each $\pi \in \Pi^\Lambda$ the matrix $A^\pi = [\sum_{i\in\pi_1} A^i, \ldots, \sum_{i\in\pi_p} A^i]$, query the C-oracle for the value $C(A^\pi)$, and pick the best. The number of operations involved and queries to the C-oracle is again $O(n^{dp^2})$. The bit size of the numbers manipulated throughout this process is polynomially bounded in the bit size of M and A, and hence the algorithm is strongly polynomial oracle time. ∎

Recall that the shaped partition polytope is defined as $\mathcal{P}^\Lambda_A = \mathrm{conv}\{A^\pi : |\pi| \in \Lambda\}$. The number of matrices in the set $\{A^\pi : |\pi| \in \Lambda\}$ is typically exponential in n, even for fixed d, p. Therefore, although the dimension of $\mathcal{P}^\Lambda_A$ is bounded by dp, this polytope potentially can have exponentially many vertices and facets as well. Lemmas

3.7 and 4.1 yield the following theorem, which shows that, in fact, shaped partition polytopes are exceptionally well behaved.

THEOREM 4.3. *Let d, p be fixed. For any $A \in \mathbb{R}^{d\times n}$ and nonempty set Λ of p-shapes of n, the number of vertices of the shaped partition polytope $\mathcal{P}_A^\Lambda$ is $O(n^{d\binom{p}{2}})$. Further, there is an algorithm that, given $n \in \mathbb{N}$, $A \in \mathbb{Q}^{d\times n}$, and oracle-presented Λ, produces all vertices of $\mathcal{P}_A^\Lambda$ in strongly polynomial oracle time using $O(n^{d^2p^3})$ operations and queries.*

Proof. By Lemma 3.7, each vertex of $\mathcal{P}_A^\Lambda$ equals the A-matrix A^π of some partition in Π_A^p. Therefore, the number of vertices of $\mathcal{P}_A^\Lambda$ is bounded above by $|\Pi_A^p|$, hence, by Lemma 4.1, is $O(n^{d\binom{p}{2}})$. To construct the set of vertices given a rational matrix A, proceed as follows. Use the algorithm of Lemma 4.1 to construct the set Π_A^p of A-generic p-partitions in strongly polynomial time using $O(n^{dp^2})$ arithmetic operations. Test the shapes of the partitions in the list to obtain its subset $\Pi^\Lambda := \{\pi \in \Pi_A^p : |\pi| \in \Lambda\}$ of A-generic Λ-partitions by querying the Λ-oracle on each of the $|\Pi_A^p| = O(n^{d\binom{p}{2}})$ partitions in Π_A^p. Construct the set of matrices $U := \{A^\pi : \pi \in \Pi^\Lambda\}$ with multiple copies identified. This set U is contained in $\mathcal{P}_A^\Lambda$, and by Lemma 3.7 contains the set of vertices of $\mathcal{P}_A^\Lambda$. So $u \in U$ will be a vertex precisely when it is not a convex combination of other elements of U. This could be tested using any linear programming algorithm, but to obtain a strongly polynomial time procedure, we proceed as follows. By Carathéodory's theorem, u will be a vertex if and only if it is not in the convex hull of any *affine basis* of $U \setminus \{u\}$. So, to test if $u \in U$ is a vertex of $\mathcal{P}_A^\Lambda$, compute the affine dimension a of $U \setminus \{u\}$. For each $(a+1)$-subset $\{u_0, \ldots, u_a\}$ of $U \setminus \{u\}$, test if it is an affine basis of $U \setminus \{u\}$, and if it is, compute the unique $\mu_0, \ldots, \mu_a$ satisfying $u = \sum_{i=0}^a \mu_i u_i$ and $\sum_{i=0}^a \mu_i = 1$. Then u is in the convex hull of $\{u_0, \ldots, u_a\}$ if and only if $\mu_0, \ldots, \mu_a \geq 0$. So u is a vertex of $\mathcal{P}_A^\Lambda$ if and only if for each affine basis we get some $\mu_i < 0$. Computing the affine dimension a, testing if an $(a+1)$-subset of $U \setminus \{u\}$ is an affine basis, and computing the μ_i can all be done by Gaussian elimination in strongly polynomial time. Since we have to perform the entire procedure for each of the $|U| \leq |\Pi^\Lambda| = O(n^{d\binom{p}{2}})$ elements $u \in U$, and for each such u the number of affine bases of $U \setminus \{u\}$ is at most $\binom{|U|-1}{dp+1}$, the number of arithmetic operations involved is $O(|U|\binom{|U|-1}{dp+1}) = O(n^{d^2p^3})$, which absorbs the work of constructing Π^Λ and obeys the claimed bound. ∎

As an immediate corollary of Theorem 4.3, we get the following polynomial bound on the number of facets of any shaped partition polytope and a strongly polynomial oracle time procedure for producing all facets (by which we mean finding, for each facet F, a hyperplane $\{X \in \mathbb{R}^{d\times p} : \langle H, X\rangle = h\}$ supporting $\mathcal{P}_A^\Lambda$ at F).

COROLLARY 4.4. *Let d, p be fixed. For any $A \in \mathbb{R}^{d\times n}$ and nonempty set Λ of p-shapes of n, the number of facets of the shaped partition polytope $\mathcal{P}_A^\Lambda$ is $O(n^{\frac{d^2p^3}{2}})$. Further, there is an algorithm that, given $n \in \mathbb{N}$, $A \in \mathbb{Q}^{d\times n}$, and oracle-presented Λ, produces all facets of $\mathcal{P}_A^\Lambda$ in strongly polynomial oracle time using $O(n^{d^2p^3})$ operations and queries.*

Proof. By the well-known upper bound theorem [18], the number of facets of any k-dimensional polytope with m vertices is $O(m^{\frac{k}{2}})$. Applying this to $\mathcal{P}_A^\Lambda$ with $k \leq dp$ and $m = O(n^{d\binom{p}{2}})$, we get the bound on the number of facets of $\mathcal{P}_A^\Lambda$. To construct the facets, first construct the set V of vertices using the algorithm of Theorem 4.3. Compute the dimension a of $\mathrm{aff}(P) = \mathrm{aff}(V)$ and compute a (possibly empty) set S of $dp - a$ points that, together with V, affinely span $\mathbb{R}^{d\times p}$. For each affinely independent

a-subset T of V, compute the hyperplane $\{X \in \mathbb{R}^{d\times p} : \langle H, X\rangle = h\}$ spanned by $S \cup T$. This hyperplane supports a facet of $\mathcal{P}_A^\Lambda$ if and only if all points in V lie on one of its closed half-spaces. Clearly, all facets of $\mathcal{P}_A^\Lambda$ are obtained that way, in strongly polynomial time and with the number of arithmetic operations and oracle queries bounded as claimed. ∎

Acknowledgment. Shmuel Onn thanks the Mathematical Sciences Research Institute at Berkeley for its support while part of this research was done.

REFERENCES

[1] N. Alon and S. Onn, *Separable partitions*, Discrete Appl. Math., 91 (1999), pp. 39–51.

[2] F. Aurenhammer and O. Schwarzkopf, *A simple on-line randomized incremental algorithm for computing higher order (Voronoi) diagrams*, in Proceedings of the 7th ACM Symposium on Computational Geometry, 1991, pp. 142–151.

[3] E.R. Barnes, A.J. Hoffman, and U.G. Rothblum, *Optimal partitions having disjoint convex and conic hulls*, Math. Programming, 54 (1992), pp. 69–86.

[4] I. Bárány and S. Onn, *Colourful linear programming and its relatives*, Math. Oper. Res., 22 (1997), pp. 550–567.

[5] A.K. Chakravarty, J.B. Orlin, and U.G. Rothblum, *Consecutive optimizers for a partitioning problem with applications to optimal inventory groupings for joint replenishment*, Oper. Res., 33 (1985), pp. 820–834.

[6] H. Edelsbrunner, P. Valtr, and E. Welzl, *Cutting dense point sets in half*, in Proceedings of the 10th ACM Symposium on Computational Geometry, 1994, pp. 203–210.

[7] P. Erdös, *The Art of Counting*, Joel Spencer, ed., MIT Press, Cambridge, MA, 1973.

[8] B. Gao, F.K. Hwang, W.-C.W. Li, and U.G. Rothblum, *Partition polytopes over* 1-*dimensional points*, 1996, Math. Progr., to appear.

[9] D. Granot and U.G. Rothblum, *The Pareto set of the partition bargaining game*, Games Econom. Behav., 3 (1991), pp. 163–182.

[10] F.K. Hwang and C.L. Mallows, *Enumerating nested and consecutive partitions*, J. Combin. Theory Ser. A, 70 (1995), pp. 323–333.

[11] F.K. Hwang, S. Onn, and U.G. Rothblum, *Representations and characterizations of the vertices of bounded-shape partition polytopes*, Linear Algebra Appl., 278 (1998), pp. 263–284.

[12] F.K. Hwang, S. Onn, and U.G. Rothblum, *Linear Programming over Partitions*, in preparation.

[13] F.K. Hwang and U.G. Rothblum, *Directional-quasi-convexity, asymmetric Schur-convexity and optimality of consecutive partitions*, Math. Oper. Res., 21 (1996), pp. 540–554.

[14] F.K. Hwang and U.G. Rothblum, *Partitions: Clustering and Optimality*, in preparation.

[15] F.K. Hwang, J. Sun, and E.Y. Yao, *Optimal set partitioning*, SIAM J. Algebraic Discrete Math., 6 (1985), pp. 163–170.

[16] P. Kleinschmidt and S. Onn, *Signable posets and partitionable simplicial complexes*, Discrete Comput. Geom., 15 (1996), pp. 443–466.

[17] L. Lovász, *On the number of halving lines*, Ann. Univ. Sci. Budapest. Eötvös Sect. Math., 14 (1970), pp. 107–108.

[18] P. McMullen, *The maximum numbers of faces of a convex polytope*, Mathematika, 17 (1970), pp. 179–184.

[19] S. Onn, *Geometry, complexity, and combinatorics of permutation polytopes*, J. Combin. Theory Ser. A, 64 (1993), pp. 31-49.

[20] A. Schrijver, *Theory of Linear and Integer Programming*, John Wiley, New York, 1986.

[21] G.M. Ziegler, *Lectures in Polytopes*, Graduate Texts in Mathematics, Springer-Verlag, New York, 1995.