Momentopes, the Complexity of Vector Partitioning, and Davenport–Schinzel Sequences*

S. Aviran¹ and S. Onn^{2,3}

¹Department of Electrical and Computer Engineering, University of California at San Diego,
La Jolla, CA 92093, USA
saviran@ucsd.edu

²Department of Operations Research, Technion - Israel Institute of Technology,
32000 Haifa, Israel
onn@ie.technion.ac.il

³Department of Mathematics, University of California at Davis,
Davis, CA 95616, USA
onn@math.ucdavis.edu

Abstract. The computational complexity of the *partition problem*, which concerns the partitioning of a 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, is determined by the number of vertices of the corresponding p -*partition polytope* defined to be the convex hull in $(d \times p)$ -space of all solutions.

In this article, introducing and using the class of *Momentopes*, we establish the lower bound $v_{p,d}(n) = \Omega(n^{\lfloor (d-1)/2 \rfloor p})$ on the maximum number of vertices of any p -partition polytope of a set of n points in d -space, which is quite compatible with the recent upper bound $v_{p,d}(n) = O(n^{d(p-1)-1})$, implying the same bound on the complexity of the partition problem.

We also discuss related problems on the realizability of Davenport–Schinzel sequences and describe some further properties of Momentopes.

1. Partition Polytopes and Partition Problems

The *partition problem* concerns the partitioning of a set A 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. More precisely, with each p -*partition* of A , namely, an ordered

* The research of Shmuel Onn was 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.

tuple $\pi = (\pi_1, \dots, \pi_p)$ of p pairwise disjoint (possibly empty) sets whose union is A , we associate the $d \times p$ matrix

$$A^\pi := \left[\sum_{a \in \pi_1} a, \dots, \sum_{a \in \pi_p} a \right] \in \mathbb{R}^{d \times p};$$

given now a convex functional $c: \mathbb{R}^{d \times p} \rightarrow \mathbb{R}$, the problem is to find a p -partition π maximizing the objective value $c(A^\pi)$ over all p -partitions of A . Such partition problems arise in a variety of areas ranging from economics to symbolic computation—see [3], [4], [6], and references therein.

The problem can be reduced to maximizing the same objective over the p -partition polytope $\mathcal{P}_A^p$ of A defined to be the convex hull in $\mathbb{R}^{d \times p}$ of all p^n matrices A^π associated with p -partitions,

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

there will always be an optimal solution which is a vertex of $\mathcal{P}_A^p$, so the problem can be solved by picking the best such vertex. When the convex functional c is presented by an evaluation oracle, it is necessary in worst case to query the oracle on every vertex. Thus, the complexity of the partition problem is intimately related to the vertex complexity $v_{p,d}(n)$ of partition polytopes, defined to be the maximal number of vertices of $\mathcal{P}_A^p$ for any set A of n points in d -space.

In [2] it was shown that if A^π is a vertex, then π is *separable*, that is, any two distinct parts can be separated by a hyperplane. The number of separable partitions was completely determined in [1] and shown to be $\Theta(n^{d \binom{p}{2}})$ for every fixed $p \geq 2$ and $d \geq 3$. This implies an upper bound on $v_{p,d}(n)$ which was exploited in [4] to provide an efficient procedure for solving the more general class of *shaped partition problems*, where partitions are restricted to being those π whose *shape* $(|\pi_1|, \dots, |\pi_p|)$ lies in a prescribed (but arbitrary) set of shapes of n .

However, as recently shown in [5], the maximal number of vertices of $\mathcal{P}_A^p$ obeys the much sharper upper bound $v_{p,d}(n) = O(n^{d(p-1)-1})$ for every fixed d and $p \geq 2$.

A main outcome of the present article is the following quite compatible lower bound.

Theorem 1.1. *For every d, p , and $n > \frac{1}{2}dp$ there is an n -set A in $\mathbb{R}^d$ such that $\mathcal{P}_A^p$ has at least $\binom{n-1}{\lfloor (d-1)/2 \rfloor p}$ vertices. Thus, for every fixed d and p , the maximal number of vertices of any p -partition polytope $\mathcal{P}_A^p$ of a set A of n points in d -space satisfies $v_{p,d}(n) = \Omega(n^{\lfloor (d-1)/2 \rfloor p})$.*

In the next section we define the Momentope $\mathcal{M}_d^p(n)$ which is a special type of partition polytope. We prove, in Theorem 2.3, a lower bound on the number of vertices of $\mathcal{M}_d^p(n)$, from which Theorem 1.1 above follows at once. In Section 3 we discuss some related problems on low degree realizability of Davenport–Schinzel sequences. We conclude with some further properties of Momentopes in Section 4.

Theorem 1.1 and the results of [5] imply that $\Omega(n^{\alpha dp}) \leq v_{p,d}(n) \leq O(n^{\beta dp})$ holds with $\beta = 1$ and any $\alpha < \frac{1}{2}$ for all sufficiently large d . It remains a challenging problem to tighten the gap $\alpha < \beta$ between the constants in these lower and upper bounds further.

2. Momentopes and the Lower Bound

Throughout, e_i stands for the i th standard unit vector in Euclidean real space. The tensor 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 $\langle U, V \rangle := \sum_{i,j} U_{i,j} \cdot V_{i,j}$. The d -moment curve is the range of the map $\mu_d: \mathbb{R} \longrightarrow \mathbb{R}^d: t \mapsto \sum_{r=1}^d t^{r-1} \cdot e_r$.

Definition 2.1 (Momentope). *The (p, d) -Momentope $\mathcal{M}_d^p(T)$ of a finite set T of real numbers is defined to be the p -partition polytope $\mathcal{P}_{\mu_d(T)}^p$ of the image $\mu_d(T)$ of T on the d -moment curve. The (p, d) -Momentope of $T = \{1, 2, \dots, n\}$ is denoted by $\mathcal{M}_d^p(n)$.*

From now on we concentrate on the Momentope $\mathcal{M}_d^p(n)$. Since its defining set $\{\mu_d(1), \dots, \mu_d(n)\}$ can be identified with $\{1, \dots, n\}$, from now a p -partition π will always mean a p -partition of $\{1, \dots, n\}$. We identify π with the function $\pi: \{1, \dots, n\} \longrightarrow \{1, \dots, p\}$ defined by $\pi(j) = i$ for all $j \in \pi_i$ and with the sequence $\pi = [\pi(1), \dots, \pi(n)]$ over $\{1, \dots, p\}$, and shall move freely among these various representations of π . We denote the $d \times p$ matrix associated with π by

$$M_d^\pi := \sum_{j=1}^n \mu_d(j) \otimes e_{\pi(j)} = \left[\sum_{j \in \pi_1} \mu_d(j), \dots, \sum_{j \in \pi_p} \mu_d(j) \right] \in \mathbb{R}^{d \times p}.$$

Thus, the Momentope $\mathcal{M}_d^p(n)$ is given by

$$\mathcal{M}_d^p(n) := \text{conv}\{M_d^\pi: \pi \text{ is a } p\text{-partition of } \{1, \dots, n\}\} \subset \mathbb{R}^{d \times p}.$$

For instance, the Momentope $\mathcal{M}_3^3(3)$ is the convex hull of the $3^3 = 27$ distinct matrices obtained by applying all possible column permutations to the following five matrices:

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 3 \\ 1 & 4 & 9 \end{bmatrix} \quad \begin{bmatrix} 0 & 2 & 1 \\ 0 & 5 & 1 \\ 0 & 13 & 1 \end{bmatrix} \quad \begin{bmatrix} 1 & 0 & 2 \\ 2 & 0 & 4 \\ 4 & 0 & 10 \end{bmatrix} \quad \begin{bmatrix} 2 & 1 & 0 \\ 3 & 3 & 0 \\ 5 & 9 & 0 \end{bmatrix} \quad \begin{bmatrix} 0 & 3 & 0 \\ 0 & 6 & 0 \\ 0 & 14 & 0 \end{bmatrix}.$$

We say that a p -partition π is d -realizable if there are real univariate polynomials $f_1, \dots, f_p$ of degree less than d such that f_i dominates on π_i , that is, for all $1 \leq i \neq k \leq p$ and all $j \in \pi_i$ we have $f_i(j) > f_k(j)$. The following lemma provides a characterization of the vertices of $\mathcal{M}_d^p(n)$.

Lemma 2.2. *The matrix M_d^π is a vertex of the Momentope $\mathcal{M}_d^p(n)$ if and only if π is d -realizable.*

Proof. Suppose first M_d^π is a vertex of $\mathcal{M}_d^p(n)$. Then there is a $d \times p$ matrix C such that the linear functional $\langle C, \cdot \rangle$ is uniquely maximized over $\mathcal{M}_d^p(n)$ at M_d^π . For $s = 1, \dots, p$ define the polynomial $f_s(t) := \sum_{r=1}^d C_{r,s} t^{r-1}$. Consider any $j \in \{1, \dots, n\}$ and let $i := \pi(j)$. Now consider any $k \neq i$ and let $\bar{\pi}$ be the partition obtained from π by

moving j from π_i to π_k . Then $M_d^\pi - M_d^{\bar{\pi}} = \mu_d(j) \otimes (e_i - e_k)$, hence

$$\begin{aligned} f_i(j) - f_k(j) &= \sum_{r=1}^d C_{r,i} j^{r-1} - \sum_{r=1}^d C_{r,k} j^{r-1} \\ &= \langle C, \mu_d(j) \otimes (e_i - e_k) \rangle = \langle C, M_d^\pi \rangle - \langle C, M_d^{\bar{\pi}} \rangle > 0. \end{aligned}$$

Thus, for all $1 \leq i \neq k \leq p$ and all $j \in \pi_i$ we have $f_i(j) > f_k(j)$ and therefore π is d -realizable.

Conversely, suppose π is d -realizable by some polynomials $f_1, \dots, f_p$, and arrange their coefficients in a $d \times p$ matrix C so that $f_s(t) = \sum_{r=1}^d C_{r,s} t^{r-1}$ for $s = 1, \dots, p$. Consider any partition $\bar{\pi} \neq \pi$. Pick any $j \in \{1, \dots, n\}$ and let $i := \pi(j)$ and $k := \bar{\pi}(j)$. We then have the inequality

$$\langle C, \mu_d(j) \otimes (e_{\pi(j)} - e_{\bar{\pi}(j)}) \rangle = \sum_{r=1}^d C_{r,i} j^{r-1} - \sum_{r=1}^d C_{r,k} j^{r-1} = f_i(j) - f_k(j) \geq 0,$$

which is strict if $i \neq k$. Since $\bar{\pi} \neq \pi$, the inequality is indeed strict for some $j \in \{1, \dots, n\}$ hence

$$\langle C, M_d^\pi \rangle - \langle C, M_d^{\bar{\pi}} \rangle = \sum_{j=1}^n \langle C, \mu_d(j) \otimes (e_{\pi(j)} - e_{\bar{\pi}(j)}) \rangle > 0.$$

It follows that $\langle C, \cdot \rangle$ is uniquely maximized over $\mathcal{M}_d^p(n)$ at M_d^π , hence M_d^π is a vertex of $\mathcal{M}_d^p(n)$. $\square$

The proof of the second part of the lemma shows, moreover, that if $\pi \neq \bar{\pi}$ are two distinct d -realizable p -partitions, then the corresponding vertices are distinct $M_d^\pi \neq M_d^{\bar{\pi}}$. For nonrealizable partitions, however, it may well happen that exponentially many π give the same matrix M_d^π .

Theorem 2.3. *For every d, p , and $n > \frac{1}{2}dp$, the Momentope $\mathcal{M}_d^p(n)$ has at least $\binom{n-1}{\lfloor (d-1)/2 \rfloor p}$ vertices.*

Proof. It suffices to prove the claim for d odd, the bound for d even being weaker. Fix then any p and $d = 2k + 1$, and consider any $n > kp$. With each choice of integers s_i^j satisfying

$$0 < s_1^1 < s_2^1 < \dots < s_p^1 < s_1^2 < s_2^2 < \dots < s_p^2 < \dots < s_1^k < s_2^k < \dots < s_p^k < n,$$

making the convention $s_p^0 := 0$ and $s_1^{k+1} := n$, we associate the p -partition defined by

$$\pi_1 := \bigcup_{j=0}^k \{s_p^j + 1, \dots, s_1^{j+1}\}, \quad \pi_i := \bigcup_{j=1}^k \{s_{i-1}^j + 1, \dots, s_i^j\}, \quad 2 \leq i \leq p.$$

We now show that π is d -realizable. Define the following polynomials of degree $d - 1 = 2k$:

$$f_1(t) := \prod_{j=1}^k (t - s_1^j - \tfrac{1}{2})(t - s_p^j - \tfrac{1}{2}), \quad f_i(t) := - \prod_{j=1}^k (t - s_{i-1}^j - \tfrac{1}{2})(t - s_i^j - \tfrac{1}{2}),$$

$$2 \leq i \leq p.$$

Each f_i flips its sign at each of its $2k$ rational roots; for $i = 1$ these are the points

$$s_1^1 + \tfrac{1}{2} < s_p^1 + \tfrac{1}{2} < s_1^2 + \tfrac{1}{2} < s_p^2 + \tfrac{1}{2} < \cdots < s_1^k + \tfrac{1}{2} < s_p^k + \tfrac{1}{2},$$

whereas for $i = 2, \dots, p$ these are the points

$$s_{i-1}^1 + \tfrac{1}{2} < s_i^1 + \tfrac{1}{2} < s_{i-1}^2 + \tfrac{1}{2} < s_i^2 + \tfrac{1}{2} < \cdots < s_{i-1}^k + \tfrac{1}{2} < s_i^k + \tfrac{1}{2}.$$

Inspecting now the sign of the leading term of f_i , which is positive for $i = 1$ and negative otherwise, it follows from the construction of π that, for each $i = 1, \dots, p$, we have that $f_i(j) > 0$ for all $j \in \pi_i$ whereas $f_i(j) < 0$ for all $j \in \{1, \dots, n\} \setminus \pi_i$. Thus, for all $1 \leq i \neq k \leq p$ and all $j \in \pi_i$ we have $f_i(j) > 0 > f_k(j)$, showing that the f_i provide a d -realization of π .

Since each of the $\binom{n-1}{kp}$ choices of s_i^j gives a different partition, and different realizable partitions π give different matrices M_d^π (see the discussion following the proof of Lemma 2.2), we find that the number of vertices of the Momentope $\mathcal{M}_d^p(n)$ is indeed at least $\binom{n-1}{kp} = \binom{n-1}{\lfloor (d-1)/2 \rfloor p}$. $\square$

Example 2.4. Let $k := p := 3$, $d := 2k + 1 = 7$, and $n := 27$, and consider the corresponding Momentope $\mathcal{M}_7^3(27)$ in $\mathbb{R}^{7 \times 3}$. Consider the following choice of s_i^j :

$$s_1^1 := 2, \quad s_2^1 := 3, \quad s_3^1 := 5, \quad s_1^2 := 9, \quad s_2^2 := 11, \quad s_3^2 := 17,$$

$$s_1^3 := 18, \quad s_2^3 := 22, \quad s_3^3 := 25;$$

the associated partition is

$$\pi := (\{1, 2, 6, 7, 8, 9, 18, 26, 27\}, \{3, 10, 11, 19, 20, 21, 22\},$$

$$\{4, 5, 12, 13, 14, 15, 16, 17, 23, 24, 25\})$$

or in its sequence form

$$\pi = [1, 1, 2, 3, 3, 1, 1, 1, 1, 2, 2, 3, 3, 3, 3, 3, 1, 2, 2, 2, 3, 3, 3, 1, 1];$$

and the polynomials providing a 7-realization of π are

$$f_1(t) := (t - \tfrac{5}{2})(t - \tfrac{11}{2})(t - \tfrac{19}{2})(t - \tfrac{35}{2})(t - \tfrac{37}{2})(t - \tfrac{51}{2}),$$

$$f_2(t) := -(t - \tfrac{5}{2})(t - \tfrac{7}{2})(t - \tfrac{19}{2})(t - \tfrac{23}{2})(t - \tfrac{37}{2})(t - \tfrac{45}{2}),$$

$$f_3(t) := -(t - \tfrac{7}{2})(t - \tfrac{11}{2})(t - \tfrac{23}{2})(t - \tfrac{35}{2})(t - \tfrac{45}{2})(t - \tfrac{51}{2}).$$

The $f_i(t)$ are plotted in Fig. 1 which shows that their domination pattern agrees with π , and hence demonstrates that the 7×3 matrix M_7^π is a vertex of the Momentope $\mathcal{M}_7^3(27)$.

Fig. 1. The 7-realization of the partition π in Example 2.4.

3. Which Davenport–Schinzel Sequences Are Realizable?

We say that the sequence $\pi = [\pi(1), \dots, \pi(n)]$ over $\{1, \dots, p\}$ is a (p, d) -sequence if for any pair $1 \leq r \neq s \leq p$, the length of any subsequence $[r, s, r, s, \dots]$ of π whose elements alternate between r and s is at most d . If, in addition, every two consecutive terms of π are distinct, then π is a *Davenport–Schinzel* (p, d) -sequence (termed so in [7]). For instance,

$$\pi = [1, 2, 1, 2, 1, 3, 1, 3, 1, 3, 2, 3, 2, 3]$$

is a Davenport–Schinzel $(3, 6)$ -sequence of maximum possible length 14.

Proposition 3.1. *If π is d -realizable (hence M_d^π a vertex of $\mathcal{M}_d^p(n)$), then it is a (p, d) -sequence.*

Proof. Suppose the polynomials $f_1, \dots, f_p$ realize π and consider any pair $1 \leq r \neq s \leq p$ and any subsequence $[\pi(j_1), \dots, \pi(j_k)]$ of π of length k whose elements alternate between r and s . Then the k values $f_{r,s}(j_1), \dots, f_{r,s}(j_k)$ of the polynomial $f_{r,s} := f_r - f_s$ alternate in sign. Since $f_{r,s}$ has degree less than d , hence at most $d - 1$ real zeros, it must be that $k \leq d$. $\square$

Is every (p, d) -sequence d -realizable? The answer is negative: there are many more (p, d) -sequences than d -realizable ones as we now show. Let $\lambda_d(p)$ denote the maximal possible length of any Davenport–Schinzel (p, d) -sequence.

Theorem 3.2. *For every fixed d and $p \geq 2$, the number of (p, d) -sequences of length n satisfies $s_{p,d}(n) = \Omega(n^{\lambda_d(p)-1})$, whereas the number of d -realizable ones satisfies $r_{p,d}(n) = O(n^{d(p-1)-1})$.*

Proof. For the estimate on $r_{p,d}(n)$, note that by Lemma 2.2 it equals the number of vertices of the Momentope $\mathcal{M}_d^p(n)$ and hence obeys the upper bound of [5] of $O(n^{d(p-1)-1})$ on the number of vertices of any partition polytope. For the estimate on $s_{p,d}(n)$, following [1]

pick any Davenport–Schinzel (p, d) -sequence $\pi = [\pi(1), \dots, \pi(k)]$ of maximal length $k := \lambda_d(p)$. Then, for all $n \geq k$, each of the $\binom{n-1}{k-1} = \Omega(n^{\lambda_d(p)-1})$ tuples $(a_1, \dots, a_k)$ of k positive integers summing to n gives a distinct (p, d) -sequence of length n by replicating a_i times the term $\pi(i)$ for each i . $\square$

Since $\lambda_d(p)$ grows superlinearly with p for any fixed $d \geq 4$ [8] (in particular, $\lambda_4(p) = \Omega(p\alpha(p))$ with $\alpha(\cdot)$ the inverse of Ackermann’s function), and $\lambda_d(p) > (d - \frac{1}{2}p)\binom{p}{2}$ for all $d \geq p$ [7], it follows that $s_{p,d}(n)$ grows much faster than $r_{p,d}(n)$ hence most (p, d) -sequences are not d -realizable.

We are led to consider a weaker notion of realization: a sequence $\pi = [\pi(1), \dots, \pi(n)]$ over $\{1, \dots, p\}$ is *weakly d -realizable* if there are reals $t_1 < \dots < t_n$ and univariate polynomials $f_1, \dots, f_p$ of degree less than d such that for all $1 \leq i \neq k \leq p$ and all $j \in \pi_i$ we have $f_i(t_j) > f_k(t_j)$. It is easy to see that a (p, d) -sequence is weakly d -realizable if and only if so is the Davenport–Schinzel (p, d) -sequence obtained from it by identifying any identical consecutive terms. While many other types of realizations of Davenport–Schinzel sequences, such as by line segments, are quite well understood (see the extensive treatment in [8]), the following interesting question remains: which Davenport–Schinzel (p, d) -sequences are weakly d -realizable?

4. More on Momentopes

We conclude with some further discussion of Momentopes. In particular, we determine the precise number of vertices for the case of $p = 2$. Let $T = \{t_1, \dots, t_n\}$ be any n -set of real numbers assumed to be increasingly ordered $t_1 < \dots < t_n$, and let $A := \mu_d(T) = \{\mu_d(t_1), \dots, \mu_d(t_n)\}$ be its image on the moment curve. As before, any p -partition π can be identified with the sequence $\pi = [\pi(1), \dots, \pi(n)]$ over $\{1, \dots, p\}$ via the obvious correspondence between A and $\{1, \dots, n\}$.

We say that a sequence $\pi = [\pi(1), \dots, \pi(n)]$ over $\{1, \dots, p\}$ is (d, T) -realizable if there are univariate polynomials $f_1, \dots, f_p$ of degree less than d such that for all $1 \leq i \neq k \leq p$ and all $j \in \pi_i$ we have $f_i(t_j) > f_k(t_j)$. Thus, with this terminology, π is d -realizable if it is (d, T) -realizable for $T = \{1, \dots, n\}$, and it is weakly d -realizable if it is (d, T) -realizable for some n -set T . Further, π will be called *strongly d -realizable* if it is (d, T) -realizable for every n -set T . The proof of Lemma 2.2 can be easily extended to the following statement for general Momentopes.

Lemma 4.1. *Let $T = \{t_1, \dots, t_n\}$ with $t_1 < \dots < t_n$ and let $A := \mu_d(T)$. Then the matrix A^π is a vertex of the Momentope $\mathcal{M}_d^p(T) = \mathcal{P}_A^p$ if and only if the sequence π is (d, T) -realizable.*

The following lemma characterizes strongly realizable sequences for $p = 2$.

Lemma 4.2. *A sequence π over $\{1, 2\}$ is a $(2, d)$ -sequence if and only if it is strongly d -realizable.*

Proof. Every strongly d -realizable π is in particular d -realizable, hence by Proposition 3.1 a (p, d) -sequence. Conversely, let π be any $(2, d)$ -sequence. Without loss of generality assume that $\pi(1) = 1$, and let $k < d$ be the number of alterations between 1 and 2. Assume $k \geq 1$, the statement being trivial if π consist of one letter only. Set $a_0 := 1$ and for $i = 1, \dots, k$ let a_i be the smallest index larger than a_{i-1} with $\pi(a_i) \neq \pi(a_{i-1})$ (i.e., where π flips from 1 to 2 or vice versa). Then, for any $T = \{t_1, \dots, t_n\}$ with $t_1 < \dots < t_n$, the polynomial $f_1(t) := \prod_{i=1}^k (\frac{1}{2}(t_{a_{i-1}} + t_{a_i}) - t)$ of degree $k < d$ and the identically zero polynomial $f_2(t) := 0$ provide a (d, T) -realization of π . $\square$

We obtain the following theorem.

Theorem 4.3. *For every $n \geq d$ and $T = \{t_1, \dots, t_n\}$ with $t_1 < \dots < t_n$, the Momentope $\mathcal{M}_d^2(T)$ has precisely $2 \sum_{k=0}^{d-1} \binom{n-1}{k}$ vertices, which are all the matrices corresponding to $(2, d)$ -sequences.*

Proof. Lemmas 4.1 and 4.2 together with an easy extension of Proposition 3.1 asserting that any (d, T) -realizable sequence π over $\{1, \dots, p\}$ is a (p, d) -sequence imply that the vertices are indeed all matrices corresponding to $(2, d)$ -sequences. For each $0 \leq k < d$, the number of two-letter sequences with precisely k alterations is $2 \binom{n-1}{k}$, yielding the stated number of vertices. $\square$

However, already for $p = 3$ and $d = 5$, the situation is more subtle and an exact expression for the number of vertices of $\mathcal{M}_d^p(n)$ is elusive: indeed, by Theorem 3.2 and the discussion following it, the number of $(3, 5)$ -sequences of length n is $\Omega(n^{10})$ whereas the number of 5-realizable ones is $O(n^9)$, so there are many non-5-realizable, hence nonstrongly 5-realizable, such sequences.

Theorem 4.3 implies that a $(2, d)$ -sequences π gives a vertex of $\mathcal{M}_d^2(T)$ regardless of the actual reals $t_1 < \dots < t_n$ forming the set T . This resembles the behavior of the classical cyclic polytope $C_d(T) = \text{conv}(\mu_d(T))$ whose combinatorial type is independent of the actual choice of its defining set on the moment curve; however, for $p \geq 3$, the existence of nonstrongly d -realizable sequences indicates that, in general, the combinatorial type of $\mathcal{M}_d^p(T)$ may depend on the defining set T .

We finish by introducing the more general class of *Shaped Momentopes* which, in particular, captures the cyclic polytopes as well. A p -shape of n is a p -tuple $(\lambda_1, \dots, \lambda_p)$ of nonnegative integers summing to n . Following [4], the *shape* of a p -partition π is the p -shape $|\pi| := (|\pi_1|, \dots, |\pi_p|)$ which counts the number of elements in each part of π . Given a set A of n points in d -space and a nonempty set Λ of p -shapes of n , the corresponding *shaped partition polytope* is defined to be the convex hull of all matrices of suitably shaped partitions,

$$P_A^\Lambda := \text{conv} \{A^\pi : \pi \text{ is a } p\text{-partition of } A, |\pi| \in \Lambda\} \subset \mathbb{R}^{d \times p}.$$

Specializing the definition to Momentopes, given a set T of n real numbers and a nonempty set Λ of p -shapes of n , the corresponding *Shaped Momentope* is defined to be $\mathcal{M}_d^\Lambda(T) := P_A^\Lambda$ with $A := \mu_d(T)$ the image of T on the d -moment curve.

The usual Momentope $\mathcal{M}_d^p(T)$ is recovered as $\mathcal{M}_d^\Lambda(T)$ with Λ the set of all p -shapes of n , and the classical cyclic polytope $C_d(T)$ is affinely equivalent to the Shaped Momentope $\mathcal{M}_d^\Lambda(T)$ for $p = 2$ and $\Lambda = \{(1, n - 1)\}$ consisting of the single 2-shape $(1, n - 1)$ of n , obtained by the simple map $\mathbb{R}^{d \times 2} \longrightarrow \mathbb{R}^d$ which erases the second column of each $d \times 2$ matrix.

We hope this article will stimulate further research on this complex and rich class of polytopes intimately related to the polynomial realization problem for Davenport–Schinzel sequences.

References

1. N. Alon and S. Onn, Separable partitions, *Discrete Applied Mathematics*, 91:39–51, 1999.
2. E. R. Barnes, A. J. Hoffman, and U. G. Rothblum, Optimal partitions having disjoint convex and conic hulls, *Mathematical Programming*, 54:69–86, 1992.
3. D. Granot and U. G. Rothblum, The Pareto set of the partition bargaining game, *Games and Economic Behavior* 3:163–182, 1991.
4. F. K. Hwang, S. Onn, and U. G. Rothblum, A polynomial time algorithm for shaped partition problems, *SIAM Journal on Optimization*, 10:70–81, 1999.
5. S. Onn and L. J. Schulman, The vector partition problem for convex objective functions, *Mathematics of Operations Research*, 26:583–590, 2001.
6. S. Onn and B. Sturmfels, Cutting corners, *Advances in Applied Mathematics*, 23:29–48, 1999.
7. D. Roselle and R. Stanton, Some properties of Davenport–Schinzel sequences, *Acta Arithmetica*, 17:355–362, 1971.
8. M. Sharir and P. K. Agarwal, *Davenport–Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

Received April 4, 2001, and in revised form October 3, 2001. Online publication February 28, 2002.