J Comb Optim (2014) 27:567–573
DOI 10.1007/s10878-012-9536-1

Are there more almost separable partitions than separable partitions?

Fei-Huang Chang · Hong-Bin Chen · Frank K. Hwang

Published online: 16 August 2012

Abstract A partition of a set of n points in d-dimensional space into p parts is called an (almost) *separable partition* if the convex hulls formed by the parts are (almost) pairwise disjoint. These two partition classes are the most encountered ones in clustering and other partition problems for high-dimensional points and their usefulness depends critically on the issue whether their numbers are small. The problem of bounding separable partitions has been well studied in the literature (Alon and Onn in Discrete Appl. Math. 91:39–51, 1999; Barnes et al. in Math. Program. 54:69–86, 1992; Harding in Proc. Edinb. Math. Soc. 15:285–289, 1967; Hwang et al. in SIAM J. Optim. 10:70–81, 1999; Hwang and Rothblum in J. Comb. Optim. 21:423–433, 2011a). In this paper, we prove that for $d \leq 2$ or $p \leq 2$, the maximum number of almost separable partitions is equal to the maximum number of separable partitions.

Keywords Partition · Separable partition · Optimal partition · Almost separable partition

1 Introduction

Let $d \geq 1$ and let $A = \{A^1, A^2, \ldots, A^n\}$ be a multi-set of n points (not necessarily distinct) in $\mathbb{R}^d$ (denoted by $A \in \mathbb{R}^{d \times n}$). For $p \geq 1$, a *p-partition* of A is an ordered

F.-H. Chang was supported in part by the National Science Council under grants NSC100-2115-M-003-006.

F.-H. Chang
Department of Mathematics and Science, National Taiwan Normal University, New Taipei, Taiwan

H.-B. Chen (✉)
Institute of Mathematics, Academia Sinica, Taipei 10617, Taiwan
e-mail: hbchen@math.sinica.edu.tw

F.K. Hwang
Department of Applied Mathematics, National Chiao-Tung University, Hsinchu 30050, Taiwan

p-tuple $\pi = (\pi_1, \ldots, \pi_p)$, where $\pi_1, \ldots, \pi_p$ are (possibly empty) pairwise disjoint subsets of A whose union is A. Note that when points are not necessarily distinct, then two partitions are considered the same if they differ only in equivalent elements. For example, if $A = \{A^1, A^2, A^3, A^4\}$ with $A^1 = A^2 = A^3 \neq A^4$, then $\pi^1 = (\pi_1^1, \pi_2^1, \pi_3^1)$ and $\pi^2 = (\pi_1^2, \pi_2^2, \pi_3^2)$ are the same partition for $\pi_1^1 = \{A^1\}$, $\pi_2^1 = \{A^2, A^4\}$, $\pi_3^1 = \{A^3\}$ and $\pi_1^2 = \{A^2\}$, $\pi_2^2 = \{A^3, A^4\}$, $\pi_3^2 = \{A^1\}$. We refer to $\pi_1, \pi_2, \ldots, \pi_p$ as the *parts* of π, to p as the *size* of π and to $(|\pi_1|, \ldots, |\pi_p|)$, where $|\pi_i|$ is the cardinality of π_i, as the *shape* of π. A typical partition problem is to find an optimal partition over a given family for a given objective function $F(\pi)$. The most general family is the constrained-shape family which requires the shape of each member partition to be in a given set Ω of shapes. This family includes four special cases which almost cover all partition families studied in the literature:

(i) A single-shape family. Ω consists of a single shape.
(ii) A bounded-shape family. For a fixed positive integer p, Ω consists of all shapes $(|\pi_1|, \ldots, |\pi_p|)$ satisfying $\sum_{j=1}^{p} |\pi_j| = n$ and $\ell_j \leq |\pi_j| \leq u_j$ for all $j = 1, \ldots, p$ where $\{\ell_j\}$ and $\{u_j\}$ are given.
(iii) A size family. Given a fixed p, it is indeed a bounded-shape family with $\ell_j = 0$ and $u_j = n$ for all $1 \leq j \leq p$.
(iv) An open family. Ω contains all shapes, i.e., without any constraint on size or shape.

If a partition problem is to maximize an objective over an X-family, then we call the problem an X problem.

In this paper we treat d and p as constants but n can be large, as is usually the case in applications. It is impractical to solve a partition problem by brute force since even for $d = 1$ and $p = 2$, the number of single-shape partitions is exponential in n as long as $\min\{|\pi_1|, |\pi_2|\} > \alpha n$ with $\alpha > 0$ a constant, not to mention for larger d and p and the other three types of problem (see the book by Hwang and Rothblum 2011b). Therefore, it is useful to identify a class of partitions known to contain an optimal partition so that we need only to search this class. A crucial condition is of course this class must be small, i.e., its cardinality is polynomial in n. Such a class is usually characterized by a partition property which every member must observe.

Barnes et al. (1992) first studied the property of "separability" with the above purpose in mind (the special cases $p = 2$ and $d = 1$ were studied by Harding (1967) and Hwang et al. (1985), respectively, earlier), which has become the most studied property in the literature of partition problems. A partition $\pi = (\pi_1, \ldots, \pi_p)$ is *separable* if for any two parts π_j and π_h of π, their convex hulls $conv(\pi_j)$ and $conv(\pi_h)$ are disjoint, i.e., $conv(\pi_j) \cap conv(\pi_h) = \emptyset$. Let $T_S^A(n, p, d)$ be the number of separable p-partitions of $A \in \mathbb{R}^d$. Call A *generic* if no $k + 1$ points of A lie in a common j-flat for any $j < k \leq d$. Define

$$H(n, d) = \sum_{j=0}^{d} \binom{n-1}{j}.$$

Harding (1967) proved $T_S^A(n, 2, d) = 2H(n, d)$, thus establishing an upper bound of $O(n^d)$ for separable 2-partitions of generic A. Hwang et al. (1999) used this upper bound to obtain an upper bound $O(n^{d\binom{p}{2}})$ for separable p-partitions, still assuming A is generic. Recently, Hwang and Rothblum (2012) extended this upper bound $O(n^{d\binom{p}{2}})$ to arbitrary A in $\mathbb{R}^d$ where A does not have to be generic and its members do not have to be distinct.

Note that the Harding's result not only yields a bound, but it is an equality. Is the equality preserved in the extension to arbitrary A? Unfortunately, it is not. To see this, call a point *multi-point* if it appears more than once in the multi-set A. Note that when A is not distinct, two partitions are considered the same (hence counted only once) if one can be obtained from the other by interchanging the same number (including 0) of copies of each multi-point. Now let A be generic and let A' consist of n copies of the same point. Then

$$T_S^{A'}(n, 2, d) = n + 1 \neq 2H(n, d) = T_S^A(n, 2, d).$$

To study the equality issue, we have to modify the Harding's result somewhat. Define

$$T_S(n, p, d) = \max_A T_S^A(n, p, d).$$

Hwang and Rothblum (2011a) proved $T_S(n, 2, d) = 2H(n, d)$.

In real world problems, there is no reason to expect the points satisfy the general position assumption, in particular, two points can be the same due to the discreteness of most measurements. When A contains multi-points, then a separable partition must have all copies of a multi-point go to the same part to preserve "disjointness". But this is too strong a requirement for a practical partition to satisfy. In particular, this requirement makes it difficult, sometimes impossible, to meet the shape constraint. Therefore, it is desirable to weaken the "strictly disjoint" condition in separable partitions to "almost disjoint" in the following sense: A partition $\pi = (\pi_1, \ldots, \pi_p)$ is called *almost separable* if for any two parts π_j and π_h of π, $conv(\pi_j) \cap conv(\pi_h) = \emptyset$ or $\{v\}$ where v in A is a vertex[1] of both $conv(\pi_j)$ and $conv(\pi_h)$.

We use the same notation of separable partitions to denote numbers of almost separable partitions by substituting the subscript S by AS. Hwang and Rothblum (2012) proved the number of almost separable partitions has the same bound as the number of separable partitions for any of shape, size or open partitions. In particular, they proved $T_{AS}(n, 2, d) = T_S(n, 2, d)$ by an algebraic approach and raised the open problem whether the equality holds for other values of p and d. In this paper we will give a geometric approach to study the equality problem. Our approach works for $d \leq 2$ or $p \leq 2$, but not for $p > 2$ and $d > 2$.

2 Main results

While our apparent goal is to prove the equality for the size family, as confirming with the literature, we will actually prove a stronger result by proving the equality for the single-shape family. Note that the equality for the single-shape family not only implies the same for the size family, but also for the constrained-shape family since

[1] A special kind of point that describes the corners of geometric shapes.

the cardinality of each of these families is simply the sum of cardinalities of all its component shapes.

Let $T_S^A(n, (n_1, \dots, n_p), d)$ denote the number of separable partitions of A with a given shape $(n_1, \dots, n_p)$, and define

$$T_S\big(n, (n_1, \dots, n_p), d\big) = \max_A T_S^A\big(n, (n_1, \dots, n_p), d\big).$$

Similarly, we define these terms for almost separable partitions by simply changing the subscript S to AS.

By definition, every separable partition is also an almost separable partition. Thus, we obtain the following results immediately.

Lemma 1 $T_S(n, (n_1, \dots, n_p), d) \le T_{AS}(n, (n_1, \dots, n_p), d)$.

We now prove $T_S(n, (n_1, \dots, n_p), d) \ge T_{AS}(n, (n_1, \dots, n_p), d)$ when $d \le 2$ or $p \le 2$, and then, together with Lemma 1, obtain the equality between the two terms.

Let $A = \{A^1, A^2, \dots, A^n\}$ be a multi-set. Suppose v is a point which appears in A $m_v > 1$ times. Then m_v is the multiplicity of v. For any $\epsilon > 0$ and any fixed point $v \in \mathbb{R}^d$, denote by $B(v, \epsilon)$ a d-dimension ball centered at the point v with radius ϵ and denote by v_ϵ an arbitrary point in $B(v, \epsilon)$. Without loss of generality, assume $A^1 = A^2 = \cdots = A^{m_v} = v$. Notice again that there is no difference among those partitions induced by interchanging $A^1, A^2, \dots, A^{m_v}$.

Lemma 2 *Suppose $\pi = (\pi_1, \dots, \pi_p)$ is an almost separable partition of $A \in \mathbb{R}^{d\times n}$ and v is a multi-point with a unique π_j containing all copies of v. Then there exists a sufficiently small real number $\epsilon_\pi > 0$ such that $\overline{\pi} = (\pi_1, \dots, \overline{\pi_j}, \dots, \pi_p)$, where $\overline{\pi_j}$ is obtained from π_j by replacing one copy A^* of v in π_j with v_{ϵ_π}, is an almost separable partition of $\overline{A} = A \cup \{v_{\epsilon_\pi}\} \setminus \{A^*\}$.*

Proof To prove the lemma, it suffices to show that there exists a sufficiently small $\epsilon_\pi > 0$ such that $conv(\overline{\pi_j}) \cap conv(\pi_h) \subseteq conv(\pi_j) \cap conv(\pi_h)$ for all $h \ne j$.

Since $\pi = (\pi_1, \dots, \pi_p)$ is an almost separable partition of A, for each π_h with $h \ne j$ we can find a separating hyperplane $H_{j,h}$ of π_h and π_j and $H_{j,h}$ does not contain the point v (this can be done by perturbing $H_{j,h}$ slightly if $H_{j,h}$ contains v). Let $\epsilon_\pi > 0$ be the minimum distance from v to all these $H_{j,h}$'s. Then it is easy to see that, for any fixed point $v_{\epsilon_\pi} \in B(v, \epsilon_\pi)$, $conv(\overline{\pi_j}) \cap conv(\pi_h) \subseteq conv(\pi_j) \cap conv(\pi_h)$ for all $h \ne j$. □

Lemma 3 *Let $\pi = (\pi_1, \dots, \pi_p)$ be an almost separable partition of $A \in \mathbb{R}^{d\times n}$ and let v belong to t parts, say $\pi_1, \pi_2, \dots, \pi_t$, with $2 \le t \le m_v$. If $d \le 2$ or $t = 2$, then there exists a sufficiently small real number $\epsilon_\pi > 0$ such that $\overline{A}$, obtained from A by replacing a copy A^* of v with v_{ϵ_π}, has an almost separable partition $\overline{\pi} = (\pi_1, \dots, \overline{\pi_j}, \dots, \pi_p)$ where $\overline{\pi_j}$ is obtained from π_j, for some $1 \le j \le t$, by replacing A^* with v_{ϵ_π}.*

Proof Since v belongs to at least two parts, we cannot simply apply Lemma 2 here. However, if we consider only one part π_j, $j \in \{1, 2, \dots, t\}$, and its separability with parts not containing v, then the proof of Lemma 2 applies to conclude the existence of a small enough $\epsilon_\pi > 0$ such that $conv(\overline{\pi_j}) \cap conv(\pi_h) \subseteq conv(\pi_j) \cap conv(\pi_h)$ for all h with π_h not containing v.

Next, we consider the separability between π_j and other parts containing v. Our argument does not apply to any such π_j, but to a specific π_j chosen in the following way. It suffices to show that for an arbitrary point $v_{\epsilon_\pi} \in B(v, \epsilon_\pi)$ there exists a part $\pi_j \in \{\pi_1, \pi_2, \ldots, \pi_t\}$ such that $conv(\overline{\pi_j}) \cap conv(\pi_h) \subseteq conv(\pi_j) \cap conv(\pi_h)$ for all $1 \le h \le t$ and $h \neq j$. This, together with $conv(\overline{\pi_j}) \cap conv(\pi_h) \subseteq conv(\pi_j) \cap conv(\pi_h)$ for all $h \notin \{1, 2, \ldots, t\}$, implies $\overline{\pi} = (\pi_1, \ldots, \overline{\pi_j}, \ldots, \pi_p)$ is an almost separable partition of $\overline{A}$.

To this aim, we claim the following and prove later: For $d \le 2$ or $t = 2$, $\mathbb{R}^d$ can be split at v into t convex spaces $\gamma_1, \ldots, \gamma_t$ with overlapping boundaries but disjoint interiors such that every $conv(\pi_h)$ is contained in the corresponding convex space γ_h for $h = 1, 2, \ldots, t$. Consequently, the point v_{ϵ_π} in $\mathbb{R}^d$ must lie in at least one of these t convex spaces, say γ_j (no matter whether it lies in the boundary or in the interior). Then $conv(\pi_j \cup \{v_{\epsilon_\pi}\}) \cap conv(\pi_h) \setminus \{v\} = \emptyset$, which implies $conv(\overline{\pi_j}) \cap conv(\pi_h) \setminus \{v\} = \emptyset$, for all $1 \le h \le t$ and $h \neq j$. Therefore the resultant partition $\overline{\pi} = (\pi_1, \ldots, \overline{\pi_j}, \ldots, \pi_p)$ is an almost separable partition of $\overline{A}$.

We now prove the above claim to complete the proof. In the case $t = 2$, one hyperplane is sufficient to separate $conv(\pi_1)$ and $conv(\pi_2)$; hence $\mathbb{R}^d$ can be split into two half spaces γ_1 and γ_2 such that $conv(\pi_h)$ is contained in the corresponding convex space γ_h for $h = 1, 2$. In the case $d \le 2$, the space considered here is either a line or a plane. For $d = 1$, then necessarily $t = 2$ and we are done. For $d = 2$, then there exists a clockwise ordering of these t convex hulls (meeting at v) such that we can completely split $\mathbb{R}^d$ into t convex spaces as desired by using separating hyperplanes between each neighbor pair of convex hulls in this ordering. This completes the proof of this lemma. $\square$

The restrictions on d and t are necessary in Lemma 3. Figure 1 demonstrates a counterexample for the case $t = d = 3$. With Lemma 3, we are now ready to prove the main result that the equality holds in Lemma 1 when $d \le 2$ or $p \le 2$.

Denote $n^*(A)$ as the number of distinct points of $A = \{A^1, A^2, \ldots, A^n\}$.

Theorem 1 *$T_S(n, (n_1, \ldots, n_p), d) = T_{AS}(n, (n_1, \ldots, n_p), d)$, when $d \le 2$ or $p \le 2$.*

Proof By Lemma 1, it suffices to prove that $T_S(n, (n_1, \ldots, n_p), d) \ge T_{AS}(n, (n_1, \ldots, n_p), d)$ when $d \le 2$ or $p \le 2$. Let $A = \{A^1, A^2, \ldots, A^n\}$ satisfy $T^A_{AS}(n, (n_1, \ldots, n_p), d) = T_{AS}(n, (n_1, \ldots, n_p), d)$ and $n^*(A)$ is maximum among all such A. If $n^*(A) = n$, then all points of A are distinct, which implies that any almost separable partition of A is also a separable partition of A. Hence, $T_S(n, (n_1, \ldots, n_p), d) \ge T^A_S(n, (n_1, \ldots, n_p), d) \ge T^A_{AS}(n, (n_1, \ldots, n_p), d) = T_{AS}(n, (n_1, \ldots, n_p), d)$, as desired.

Suppose that $n^*(A) < n$, which implies the existence of a multi-point v with multiplicity $m_v > 1$. For each almost separable partition $\pi = (\pi_1, \ldots, \pi_p)$ of A with shape $(n_1, \ldots, n_p)$, there are only two cases: all copies of v belong to either exactly one part π_j for some j, or to t parts with $t \ge 2$. By Lemmas 2 and 3, for each

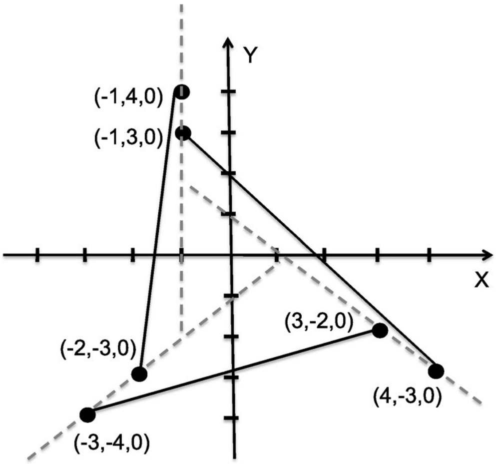

Fig. 1 Consider the multi-set $A = \{(0,0,1),(0,0,1),(0,0,1),(-1,4,0),(-2,-3,0),(-3,-4,0),(3,-2,0),(-1,3,0),(4,-3,0)\}$. Let $\pi_1 = \{(0,0,1),(-1,4,0),(-2,-3,0)\}$, $\pi_2 = \{(0,0,1),(-3,-4,0),(3,-2,0)\}$, $\pi_3 = \{(0,0,1),(-1,3,0),(4,-3,0)\}$. This figure shows only the plane $Z = 0$ for convenience. Obviously, $\pi = (\pi_1, \pi_2, \pi_3)$ is an almost separable partition. But one cannot find a partition of $\mathbb{R}^3$ into 3 convex subspaces each containing a part exclusively, as described in the proof of Lemma 3, since no γ_j, $j = 1, 2, 3$, can move into the center area of the cone-like shape (encircled by π_1, π_2 and π_3)

almost separable partition π of A we can find an almost separable partition $\overline{\pi}$ of $\overline{A} = A \cup \{v_{\epsilon_\pi}\} \setminus A^*$, where A^* is a copy of v, for sufficiently small ϵ_π. Moreover, the shape is preserved. Since there are finitely many almost separable partitions of A, we can choose the smallest ϵ_π over all almost separable partitions to be ϵ. Accordingly, for every almost separable partition π of $A = \{A^1, A^2, \ldots, A^n\}$ there exists a corresponding $\overline{\pi}$ which is an almost separable partition of $\overline{A} = A \cup \{v_\epsilon\} \setminus A^*$ and the shape is preserved. Finally, the corresponding partitions of $\overline{A}$ must be all distinct, thus implies a one-to-one correspondence. To see this, suppose that two distinct almost separable partitions $\pi^1 = (\pi_1^1, \ldots, \pi_p^1)$ and $\pi^2 = (\pi_1^2, \ldots, \pi_p^2)$ of A map to two almost separable partitions $\overline{\pi^1} = (\pi_1^1, \ldots, \overline{\pi_r^1}, \ldots, \pi_p^1)$ and $\overline{\pi^2} = (\pi_1^2, \ldots, \overline{\pi_s^2}, \ldots, \pi_p^2)$ of $\overline{A}$ and $\overline{\pi^1} = \overline{\pi^2}$. Then we have $r = s$ since $\overline{\pi_r^1}$ and $\overline{\pi_s^2}$ are the only two parts containing v_ϵ which is unique. Further, $\pi_h^1 = \pi_h^2$ for all $h \neq r$. Obviously, if we identify v_ϵ with A^*, then $\pi_r^1 = \pi_s^2$ because of $\overline{\pi_r^1} = \overline{\pi_s^2}$ (notice that v_ϵ and A^* are the only difference between π and $\overline{\pi}$). It follows $\pi^1 = \pi^2$, a contradiction to the assumption that they are distinct. As a result, there is a one-to-one correspondence from almost separable partitions of A to almost separable partitions of $\overline{A}$. Then $T_{AS}^{\overline{A}}(n, (n_1, \ldots, n_p), d) \geq T_{AS}^{A}(n, (n_1, \ldots, n_p), d)$ with $n^*(\overline{A}) = n^*(A) + 1 > n^*(A)$, a contradiction to the assumption that $n^*(A)$ is maximum. Thus $n^*(A) = n$ and Theorem 1 is proved. $\square$

Corollary 1 *$T_S(n, p, d) = T_{AS}(n, p, d)$ for $d \leq 2$ or $p \leq 2$.*

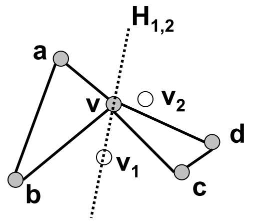

Fig. 2 This figure shows a partition $\pi = (\pi_1, \pi_2)$ with a separating line $H_{1,2}$, where $\pi_1 = \{a, b, v\}$ and $\pi_2 = \{c, d, v\}$ shares a multi-point v with two copies. The boundary of the two corresponding convex spaces is $H_{1,2}$. Suppose $v_1 = v_\epsilon$ lies on the separating line $H_{1,2}$ and is assigned to the π_1 side, as determined in Lemma 3. Then π is mapped to $\overline{\pi} = (\{a, b, v_1\}, \{c, d, v\})$. In this case the partition $\overline{\pi}' = (\{a, b, v\}, \{c, d, v_1\})$, which is almost separable, is left unmapped. Instead, suppose $v_2 = v_\epsilon$ lies on the π_2 side. Then π is mapped to $\overline{\pi} = (\{a, b, v\}, \{c, d, v_2\})$ while the unmapped partition $\overline{\pi}' = (\{a, b, v_2\}, \{c, d, v\})$ violates the almost-separability property

Remarks In the proof of Theorem 1, for any multi-point v, there exists a one to one mapping from the set of almost separable partitions on A to the set of almost separable partitions on $\overline{A}$ which differs from A by replacing a copy A^* of v with v_ϵ. An interesting question is whether this mapping is bijective, i.e., whether a member of $\overline{A}$ is left unmapped. The example given in Fig. 2 gives a positive answer and also shows that the solution depends greatly on the position where v_ϵ is placed. Actually, for a specific partition we can put the point $v_\epsilon \in B(v, \epsilon)$ in the boundary of two (out of t) convex spaces as described in Lemma 3. If v_ϵ is assigned to the left space, we obtain an almost separable partition; if to the right space, we obtain another. Since we can choose only one such space to place v_ϵ, one of the two almost separable partitions is left unmapped (see Fig. 2 as an example). As a result, we can conclude that a set A satisfying $T^A_{AS}(n, (n_1, \ldots, n_p), d) = T_{AS}(n, (n_1, \ldots, n_p), d)$ contain no multi-points when $d = 2$ and when $p = 2$ and $d \neq 1$. The condition $d \neq 1$ is necessary for otherwise the two 1-dimensional convex spaces must be two intervals (allowing the degenerating interval $\{v\}$) meeting exactly at v, such that v_ϵ must lie in only one such interval. In other words, maximum of $T_{AS}(n, (n_1, \ldots, n_p), 2)$ is always achieved by a set A consisting of distinct A^i's. Hence all elements contributing to $T_{AS}(n, (n_1, \ldots, n_p), 2)$ are separable partitions.

References

Alon N, Onn S (1999) Separable partitions. Discrete Appl Math 91:39–51

Barnes ER, Hoffman AJ, Rothblum UG (1992) Optimal partitions having disjoint convex and conic hulls. Math Program 54:69–86

Harding EF (1967) The number of partitions of a set of n points in k dimensions induced by hyperplanes. Proc Edinb Math Soc 15:285–289

Hwang FK, Rothblum UG (2011a) On the number of separable partitions. J Comb Optim 21:423–433

Hwang FK, Rothblum UG (2011b) Partitions: optimality and clustering. World Scientific, Singapore

Hwang FK, Rothblum UG (2012) Bounding the number of almost separable partitions (to appear)

Hwang FK, Sun J, Yao EY (1985) Optimal set partitioning. SIAM J Algebr Discrete Methods 6:163–170

Hwang FK, Onn S, Rothblum UG (1999) A polynomial time algorithm for shaped partition problems. SIAM J Optim 10:70–81