Approximating the Colorful Carathéodory Theorem

Wolfgang Mulzer* Yannik Stein*

Abstract

Let $P_1, \ldots, P_{d+1} \subset \mathbb{R}^d$ be d-dimensional point sets such that the convex hull of each P_i contains the origin. We call the sets P_i *color classes*, and we think of the points in P_i as having color i. A *colorful choice* is a set with at most one point from each color class. The *colorful Carathéodory theorem* guarantees the existence of a colorful choice whose convex hull contains the origin. So far, the computational complexity of finding such a colorful choice is unknown.

An *m-colorful choice* is a set that contains at most m points from each color class. We present an approximation algorithm that computes for any constant $\varepsilon > 0$, an $\lceil \varepsilon(d+1) \rceil$-colorful choice containing the origin in its convex hull in polynomial time. This notion of approximation has not been studied before, and it is motivated through the applications of the colorful Carathéodory theorem in the literature. Second, we show that the exact problem can be solved in $d^{O(\log d)}$ time if $\Theta(d^2 \log d)$ color classes are available, improving over the trivial $d^{O(d)}$ time algorithm.

1 Introduction

Let $P \subset \mathbb{R}^d$ be a point set. Carathéodory's theorem [4, Theorem 1.2.3] states that if $\vec{0} \in \operatorname{conv}(P)$, there is a subset $P' \subseteq P$ of size $d+1$ with $\vec{0} \in \operatorname{conv}(P')$. Bárány [1] gives a *colorful* generalization.

Theorem 1 (Colorful Carathéodory Theorem) *Let $P_1, \ldots, P_{d+1} \subset \mathbb{R}^d$ be point sets (the color classes) with $\vec{0} \in \operatorname{conv}(P_i)$, for $i = 1, \ldots, d+1$. There is a* colorful choice *C with $\vec{0} \in \operatorname{conv}(C)$. Here, a colorful choice is a set with at most one point from each color class.*

Theorem 1 yields Carathéodory's theorem by setting $P_1 = \cdots = P_{d+1}$. Moreover, there are many variants with weaker assumptions [5]. While Carathéodory's theorem has a proof that gives a polynomial-time algorithm, very little is known about the algorithmic complexity of the colorful Carathéodory theorem [2]. This question is particularly interesting since Sarkaria's proof [10] of Tverberg's theorem [11] can be interpreted as a polynomial-time reduction from computing Tverberg partitions to computing a colorful choice with the origin in its convex hull. Both problems lie in *Total Function NP* (TFNP), the complexity class of total search problems that are solvable in non-deterministic polynomial time. It is well known that no problem in TFNP is NP-hard unless NP = coNP [3].

Related problems have been shown to be complete for subclasses of TFNP. Recently, Meunier and Sarrabezolles [6] proved that given $d+1$ pairs of points $P_1, \ldots, P_{d+1} \in \mathbb{Q}^d$ and a colorful choice that contains the origin in its convex hull, it is PPAD-complete [9] to find another colorful choice with the origin in its convex hull. The authors [8] showed the following generalization of the colorful Carathéodory problem to be PLS-complete [3]: given sets $P_1, \ldots, P_n \subset \mathbb{R}^d$, find a colorful choice s.t. the distance of its convex hull to the origin cannot be decreased by swapping a *single* point with another point of the same color.

Since we have no polynomial-time algorithms for the colorful Carathéodory theorem, approximation algorithms are of interest. This was first studied by Bárány and Onn [2] who described how to find a colorful choice whose convex hull is "close" to the origin. Let $\varepsilon, \rho > 0$ be parameters. Given sets $P_1, \ldots, P_{d+1} \in \mathbb{Q}^d$ encoded in L bits s.t. (i) each P_i contains a ball of radius ρ centered at the origin in its convex hull; and (ii) all points $p \in P_i$ fulfill $1 \leq \|p\| \leq 2$, one can find a colorful choice C s.t. $d(\vec{0}, \operatorname{conv}(C)) \leq \varepsilon$ in time $\operatorname{poly}(L, \log(\varepsilon^{-1}), \rho^{-1})$ on the WORD-RAM with logarithmic costs. If $\rho^{-1} = L^{O(1)}$, the algorithm actually guarantees $\vec{0} \in \operatorname{conv}(C)$.

However, when using the colorful Carathéodory theorem in a proof, it is often crucial that the colorful choice contains the origin in its convex hull. Being "close" is not enough. On the other hand, allowing multiple points from each color class may have a natural interpretation in the reduction. This is the case in Sarkaria's proof [10] of Tverberg's theorem and in the proof of the First Selection Lemma [4, Theorem 9.1.1]. This motivates a different notion of approximation. Given a parameter m and sets $P_1, \ldots, P_{d+1} \in \mathbb{Q}^d$, find a set C s.t. $\vec{0} \in \operatorname{conv}(C)$ and s.t. $|C \cap P_i| \leq m$ for all P_i. In contrast to Bárány and Onn's setting, we have no general position assumption. Surprisingly, this notion does not seem to have been studied before.

*Institut für Informatik, Freie Universität Berlin, Germany. {mulzer,yannikstein}@inf.fu-berlin.de. W. Mulzer is partially supported by DFG Grants MU 3501/1 and MU 3501/2. Y. Stein is supported by the Deutsche Forschungsgemeinschaft within the research training group 'Methods for Discrete Structures' (GRK 1408).

Our Results. Given sets $P_1, \ldots, P_n \subset \mathbb{R}^d$, we call a set C containing at most m points from each set P_i an *m-colorful choice*. A 1-colorful choice is also called *perfect colorful choice*. All presented algorithms are analyzed on the Real-Ram model with unit costs. We begin with an algorithm based on a dimension reduction argument that repeatedly combines approximations for lower dimensional linear subspaces. This leads to the following result:

Theorem 2 *Let $P_1, \ldots, P_{d+1} \subset \mathbb{R}^d$ be sets of size at most $d+1$ s.t. $\vec{0} \in \mathrm{conv}(P_i)$ for $i = 1, \ldots, d+1$. Then, for any $\varepsilon = \Omega(d^{-1/3})$, an $\lceil \varepsilon(d+1) \rceil$-colorful choice containing the origin in its convex hull can be computed in $d^{O((1/\varepsilon)\log(1/\varepsilon))}$ time.*

In particular, for any constant ε the algorithm from Theorem 2 runs in polynomial-time. Given $\Theta(d^2 \log d)$ color classes, we can also improve the naive $d^{O(d)}$ algorithm for finding a perfect colorful choice.

Theorem 3 *Let $P_1, \ldots, P_n \subset \mathbb{R}^d$ be $n = \Theta(d^2 \log d)$ sets of size at most $d+1$ s.t. $\vec{0} \in \mathrm{conv}(P_i)$, for $i = 1, \ldots, n$. Then, a perfect colorful choice can be computed in $d^{O(\log d)}$ time.*

2 Fundamentals

Throughout the paper, we denote for a given point set $P = \{p_1, \ldots, p_n\} \subset \mathbb{R}^d$ by $\mathrm{span}(P) = \{\sum_{i=1}^n \alpha_i p_i \mid \alpha_i \in \mathbb{R}\}$ its linear span and by $\mathrm{span}(P)^\perp = \{v \in \mathbb{R}^d \mid \forall p \in \mathrm{span}(P) : \langle v, p \rangle = 0\}$ the subspace orthogonal to $\mathrm{span}(P)$; by $\mathrm{aff}(P) = \{\sum_{i=1}^n \alpha_i p_i \mid \alpha_i \in \mathbb{R}, \sum_{i=1}^n \alpha_i = 1\}$ its affine hull; by $\mathrm{pos}(P) = \{\sum_{i=1}^n \mu_i p_i \mid \mu_i \geq 0\}$ all linear combinations with nonnegative coefficients; by $\mathrm{conv}(P) = \{\sum_{i=1}^n \lambda_i p_i \mid \lambda_i \geq 0, \sum_{i=1}^n \lambda_i = 1\}$ its convex hull; and by $\dim(P)$ the dimension of $\mathrm{span}(P)$.

Furthermore, we say that a set $P \subset \mathbb{R}^d$ is in *general position* if for every $k \leq d$, no $k+2$ points lie in a k-flat *and* if no proper subset of P contains the origin in its convex hull. We also use the following constructive version of Carathéodory's theorem:

Lemma 4 *Let $P \subset \mathbb{R}^d$ be a set of $O(d)$ points s.t. $\vec{0} \in \mathrm{conv}(P)$. In $O(d^4)$ time, we can find a subset $P' \subseteq P$ of at most $d+1$ points in general position such that P' contains the origin in its convex hull.* ☐

3 Approximation by Rebalancing

Let $P_1, \ldots, P_{d+1} \subset \mathbb{R}^d$ be the color classes and $\lceil \varepsilon(d+1) \rceil$ be the desired approximation guarantee. Throughout the algorithm, we maintain a temporary approximation $C \subset P_1 \cup \cdots \cup P_{d+1}$ that contains the origin in its convex hull, but may have more than $\lceil \varepsilon(d+1) \rceil$ points of the same color. The algorithm then repeatedly replaces at least one point from each color class that appears more than $\lceil \varepsilon(d+1) \rceil$ times in C by colors that appear only "few" times using a dimension reduction argument.

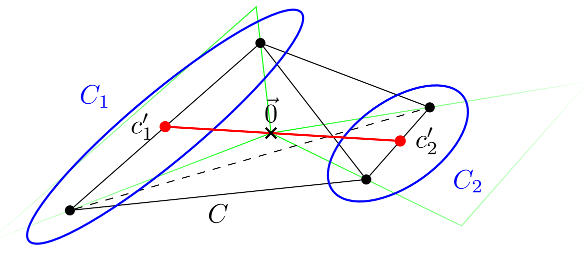

Figure 1: Example of Lemma 6 in three dimensions.

The following lemma enables us to replace a single point in C by an approximate colorful choice for the orthogonal space $\mathrm{span}(C)^\perp$:

Lemma 5 *Let $C \subset \mathbb{R}^d$, $|C| = k \leq d+1$, be a set in general position that contains the origin in its convex hull. Furthermore, let $Q \subset \mathbb{R}^d$ be a set of size $O(d)$ whose orthogonal projection onto $\mathrm{span}(C)^\perp$ contains the origin in its convex hull. Then, there is a point $c \in C$ computable in $O(d^4)$ time s.t. $\vec{0} \in \mathrm{conv}(Q \cup C \setminus \{c\})$.*

Proof. Write $Q = \{q_1, \ldots, q_l\}$. Each q_i can be expressed as $\widetilde{q}_i + \hat{c}_i$, where $\widetilde{q}_i$ denotes the orthogonal projection of q_i onto $\mathrm{span}(C)^\perp$ and $\hat{c}_i \in \mathrm{span}(C)$. By our assumption, the origin is a convex combination of $\widetilde{q}_1, \ldots, \widetilde{q}_l$: $\vec{0} = \sum_{i=1}^l \lambda_i \widetilde{q}_i$, where $\lambda_i \geq 0$ and $\sum_{i=1}^l \lambda_i = 1$. Consider the convex combination $q = \sum_{i=1}^l \lambda_i q_i$ of points in Q with the same coefficients. Since $q = \sum_{i=1}^l \lambda_i q_i = \sum_{i=1}^l \lambda_i(\widetilde{q}_i + \hat{c}_i) = \sum_{i=1}^l \lambda_i \hat{c}_i$, q is contained in $\mathrm{span}(C)$. By our assumption, we have $\vec{0} \in \mathrm{conv}(C)$ and C is in general position. It can be easily verified that this implies $\mathrm{pos}(P) = \mathrm{span}(C)$. Thus, there are $k-1$ points $c_{j_1}, \ldots, c_{j_{k-1}}$ in C s.t. $-q \in \mathrm{pos}(c_{j_1}, \ldots, c_{j_{k-1}})$. We take $c \in C$ as the single point that does not appear in $c_{j_1}, \ldots, c_{j_{k-1}}$. It can be found in $O(d^4)$ time by solving $k \leq d+1$ linear systems of equations $L_1, \ldots, L_k$, where L_j is defined as $\sum_{c_i \in C, i \neq j} \alpha_i c_i = -q$. Since C is in general position, all $(k-1)$-subsets of C are a basis for $\mathrm{span}(C)$. Thus, the linear systems have unique solutions. Since $\vec{0} \in \mathrm{conv}(C)$, one of the linear systems has a solution with no negative coefficients. ☐

Unfortunately, we do not know how to influence the color of c in Lemma 5. We would like to replace a point whose color contributes more than $\lceil \varepsilon(d+1) \rceil$ points to C. The next lemma gives us more control.

Lemma 6 *Let $C \subset \mathbb{R}^d$, $|C| \leq d+1$, be a set in general position s.t. $\vec{0} \in \mathrm{conv}(C)$ and let $C_1, \ldots, C_m$ be a partition of C. Then, we can find in $O(d^3)$ time a set $C' = \{c_1', \ldots, c_m'\} \subset \mathbb{R}^d$ with the following properties:*

(1) $\forall i = 1, \dots, m$: $c'_i \in \mathrm{pos}(C_i) \setminus \{\vec{0}\}$; *(2)* $\vec{0} \in \mathrm{conv}(C')$; *and (3)* $\dim(C') = m - 1$. *We call the points in* C' representatives *for* C *w.r.t. the partition* $C_1, \dots, C_m$.

Proof. Since C contains the origin in its convex hull, we can write $\vec{0}$ as $\vec{0} = \sum_{c \in C} \lambda_c c$, where all $\lambda_c > 0$, since C is in general position. Define c'_j as $c'_j = \sum_{c \in C_j} \lambda_c c$ for $i = 1, \dots, m$. Properties 1. and 2. can be easily verified for the set $C' = \{c'_1, \dots, c'_m\}$. Furthermore, c'_1 can be expressed as a linear combination of the other points in C': $c'_1 = -(c'_2 + \dots + c'_m)$. Thus, $\dim(C') < m$. On the other hand, we have $\dim(C') \geq m - 1$ due to general position. This proves Property 3. See Figure 1 for an example. ◻

Instead of applying Lemma 5 to C directly, we use Lemma 6 to obtain a carefully chosen set of representative points and apply Lemma 5 to replace a representative. By choosing the partition for the representatives appropriately, we can influence the color of the removed points.

Now, we are ready to put everything together. The algorithm repeatedly replaces points in C by a recursively computed approximate colorful choice for a linear subspace. We are given as input the color classes $P_1, \dots, P_{d+1} \subset \mathbb{R}^d$, each containing the origin in its convex hull, and the current recursion depth j. Define $\mathcal{M}(j)$ as $\mathcal{M}(j) = \lceil (1-\varepsilon)^{-j/2} \varepsilon (d+1) \rceil$ and $\mathcal{D}(j)$ as $\mathcal{D}(j) = \lceil (1-\varepsilon)^j \varepsilon (d+1) \rceil$. In recursion level j, the input is $\mathcal{D}(j)$-dimensional and the algorithm computes an $\mathcal{M}(j)$-colorful choice. Hence, in the topmost recursion level (i.e., $j = 0$), a $\lceil \varepsilon (d+1) \rceil$-colorful choice is computed. If $d = O(1)$, we compute an approximation by brute force. Otherwise, we initialize the temporary approximation C with a complete color class and prune it with Lemma 4. If the pruned set C is an $\mathcal{M}(j)$-colorful choice, we return it. Otherwise, we repeat the following balancing-steps: we partition C into $k = \mathcal{D}(j) - \mathcal{D}(j+1) + 1$ sets $C_1, \dots, C_k$, where the points from each color in C are distributed evenly among the k sets. Let $n_i = |P_i \cap C|$ denote the number of points from P_i in C. Since $k \leq \mathcal{M}(j) + 1$, each set in the partition contains at least one point from each color class P_i for which $n_i \geq \mathcal{M}(j) + 1$. Applying Lemma 6, we compute representatives $C' = \{c'_1, \dots, c'_k\}$ for this partition. Note that $\dim(C') = k - 1$ and that $\dim(\mathrm{span}(C')^\perp) = \mathcal{D}(j) - k + 1 = \mathcal{D}(j+1)$. We call a color class P_i *light* if $n_i \leq \mathcal{M}(j) - \mathcal{M}(j+1)$, and *heavy*, otherwise. We find $d - k + 2$ light color classes and project them orthogonally onto $\mathrm{span}(C')^\perp$. Let $\widetilde{P}_{j_1}, \dots, \widetilde{P}_{j_{d-k+2}}$ denote the projections. Next, we recursively compute an $\mathcal{M}(j+1)$-colorful choice $\widetilde{Q}$ for the space orthogonal to $\mathrm{span}(C')$ with $(\widetilde{P}_{j_1}, \dots, \widetilde{P}_{j_{\mathcal{D}(j+1)+1}}, j+1)$ as input. Let Q be the point set whose projection gives $\widetilde{Q}$. Using Lemma 5, we compute a point $c'_j \in C'$ s.t. $\mathrm{conv}(Q \cup C' \setminus c'_j)$ contains the origin. We replace the subset C_j of C by Q and prune C again with Lemma 4. Since each representative c'_i is contained in the cone $\mathrm{pos}(C_i)$, $Q \cup C \setminus C_j$ still contains the origin in its convex hull and the invariant is maintained. Thus, in each iteration, at least one point from each color class P_i for which $n_i > \mathcal{M}(j)$ is replaced by points from light color classes. This is repeated until no color class appears more than $\mathcal{M}(j)$ times in C.

Proof of Theorem 2. We prove correctness by induction on the recursion depth. In particular, we show that the input in the jth recursion is $\mathcal{D}(j)$-dimensional and that an $\mathcal{M}(j)$-colorful choice is returned. There are two base cases: if $d = O(1)$ we compute a perfect colorful choice by brute force in $O(1)$ time. This is always a valid approximation regardless of $\mathcal{M}$. If $\mathcal{D}(j) + 1 \leq \mathcal{M}(j)$, we obtain a valid approximation by pruning C with Lemma 4. Hence, the claim holds in both base cases. In each level of the recursion, the dimension is reduced by $k - 1$. The dimension of the input in the recursion is thus $\mathcal{D}(j) - k - 1 = \mathcal{D}(j+1)$ as claimed. Since $\mathcal{D}$ is decreasing, some base case is reached eventually. Assume now that the current recursion depth is j and that the claim holds for all $j' > j$. Let $C^{(t)}$ denote the set C after t iterations of the balancing-steps in the jth recursion and $n_i^{(t)}$ the number of points from P_i in C. Define the *excess of a color* P_i as $e_i^{(t)} = \max\{0, n_i^{(t)} - \mathcal{M}(j)\}$ and the *excess of* $C^{(t)}$ as $\mathcal{E}(C^{(t)}) = \max_{i=1}^{d+1} e_i^{(t)}$. We show the following invariant: (α) $\vec{0} \in \mathrm{conv}(C^{(t)})$; and ($\beta$) $\mathcal{E}(C^{(t)}) < \mathcal{E}(C^{(t-1)})$ for $t \geq 1$. The invariant implies that eventually an $\mathcal{M}(j)$-colorful choice is returned.

Before the first iteration, the invariant holds since $C^{(0)} = P_1$. Assume we are now in iteration t and the invariant holds for all previous iterations. Due to Lemmas 5 and 6, we have $\vec{0} \in \mathrm{conv}(C^{(t)})$ and thus Property (α) holds. Because $C^{(t-1)}$ was not an $\mathcal{M}(j)$-colorful choice (otherwise the balancing-steps would not haven executed), $\mathcal{E}(C^{(t-1)}) \geq 1$. Since Q contains only light color classes, adding Q to $C^{(t-1)}$ does not increase the excess. At least one point in C from each color P_i with $e_i^{(t-1)} \geq 1$ is replaced by Q. Hence, $\mathcal{E}(C^{(t)}) < \mathcal{E}(C^{(t-1)})$. Finally, we show that there are always $\mathcal{D}(j+1)+1$ light color classes for the recursion. By induction, the recursively computed set Q is an $\mathcal{M}(j+1)$-colorful choice. As C is pruned to at most $\mathcal{D}(j)+1$ points at the end of the balancing-steps, there are at most $\left\lfloor \frac{\mathcal{D}(j)+1}{\mathcal{M}(j) - \mathcal{M}(j+1)} \right\rfloor$ heavy color classes. One can show that this is at most $\mathcal{D}(j) - (\mathcal{D}(j+1) + 1)$ for $d = \Omega(1/\varepsilon^3)$. Since we assumed $\varepsilon = \Omega(d^{-1/3})$, we can always find $\mathcal{D}(j+1) + 1$ light colors.

We now analyze the running time. Each iteration of the balancing-steps reduces the excess by at least one until the desired approximation guarantee is reached. Thus, the total number of iterations is bounded by $\mathcal{D}(j) + 1 - \mathcal{M}(j) = O(d)$. Each iteration requires $O(d^4)$ time. The recursion stops when $d = O(1)$ or

$\mathcal{M}(j) \geq \mathcal{D}(j) + 1$. In the first case, a perfect colorful choice is computed in $O(1)$ time. In the second case, we spend $O(d^4)$ time since pruning P_1 with Lemma 4 already gives a valid approximation. We can bound the recursion depth j until the second base case is reached. Since $\mathcal{M}(j) \geq \varepsilon(1-\varepsilon)^{j/2}(d+1)$ and $3(1-\varepsilon)^j(d+1) \geq \mathcal{D}(j)+1$, we have $\mathcal{M}(j) \geq \mathcal{D}(j)+1$ for $j = O((1/\varepsilon)\log(1/\varepsilon))$ using that $-\log(1-\varepsilon) = \Omega(\varepsilon)$. Thus, the total running time is $d^{O((1/\varepsilon)\log(1/\varepsilon))}$. ☐

4 A Subexponential Exact Algorithm

Now, we consider the case that we have “many” color classes instead of “only” $d+1$: given $\Theta(d^2 \log d)$ color classes, our algorithm computes a perfect colorful choice in $d^{O(\log d)}$ time, improving the brute force $d^{O(d)}$ algorithm. The algorithm follows the structure of Miller and Sheehy’s algorithm for computing approximate Tverberg partitions [7]. We repeatedly combine m-colorful choices (for some m) to one $\lceil m/2 \rceil$-colorful choice. Eventually, we obtain a perfect colorful choice.

Lemma 7 *Let $C_1, \ldots, C_{d+1} \subset \mathbb{R}^d$ be m-colorful choices s.t. $|C_i| \leq d+1$ and s.t. $\vec{0} \in \mathrm{conv}(C_i)$. Then, a $\lceil m/2 \rceil$-colorful choice containing the origin in its convex hull can be computed in $O(d^5)$ time.*

Proof. First, we prune each set C_i with Lemma 4. This requires $O(d^5)$ time. Next, for $i = 1, \ldots, d+1$, we partition the colorful choice C_i into two sets $C_{i,1}, C_{i,2}$ of equal size s.t. the points from each color class are distributed evenly among the two sets. For each partition $C_{i,1}, C_{i,2}$, we apply Lemma 6 to obtain two representatives $c_{i,1}$ and $c_{i,2}$ in $O(d^3)$ time. By Lemma 6, we have $c_{i,1} \in \mathrm{pos}(C_{i,1})$ and $c_{i,2} \in \mathrm{pos}(C_{i,2})$. Since $\vec{0} \in \mathrm{conv}(\{c_{i,1}, c_{i,2}\})$, both points lie on a line through the origin and thus $-c_{i,1} \in \mathrm{pos}(C_{i,2})$. The $d+1$ points $c_{1,1}, c_{2,1}, \ldots, c_{d+1,1}$ are linearly dependent, so there exists a nontrivial linear combination $\vec{0} = \alpha_1 c_{1,1} + \cdots + \alpha_{d+1} c_{d+1,1}$. For $i = 1, \ldots d+1$, we let the set C contain $C_{i,1}$ if $\alpha_i > 0$ (since $c_{i,1} \in \mathrm{pos}(C_{i,1})$) and $C_{i,2}$ if $\alpha_i < 0$ (since $-c_{i,1} \in \mathrm{pos}(C_{i,2})$). By construction, C contains the origin in its convex hull and exactly one of $C_{i,1}, C_{i,2}$, for $i = 1, \ldots, d+1$. Since all sets C_i are m-colorful choices, C is a $\lceil m/2 \rceil$-colorful choice. ☐

Proof of Theorem 3. Let A be an array of size $k = \Theta(\log d)$. We set $c_0 = d+1$ and $c_i = \lceil c_{i-1}/2 \rceil$, for $i = 1, \ldots, k-1$. The ith cell of A stores a collection of c_i-colorful choices, such that each color class appears in exactly one colorful choice in A. Initially, $A[0]$ contains all $\Theta(d^2 \log d)$ color classes. We repeat the following steps, until we have computed a perfect colorful choice: let i be the maximum index s.t. $A[i]$ contains some $d+1$ sets $C_1, \ldots, C_{d+1}$. We apply Lemma 7 to obtain one c_{i+1}-colorful choice C. Let C' be the set C pruned with Lemma 4. If C' is a perfect colorful choice, we return it. Otherwise, we add it to $A[i+1]$. Furthermore, we add all colors that were removed during the pruning to $A[0]$. As these colors do not appear anywhere else in A, the invariant is maintained. We claim that a combination of $d+1$ sets in $A[k]$ for $k = \lceil \log(d+1) \rceil + 1$ results in a perfect colorful choice. We have $c_j \leq \frac{d+1}{2^k} + 2$. Thus, sets in $A[\lceil \log(d+1) \rceil]$ are 3-colorful choices, sets in $A[\lceil \log(d+1) \rceil + 1] = A[k]$ are 2-colorful choices and the combination of $d+1$ sets in $A[k]$ gives a perfect colorful choice. It remains to show that we can always make progress. The array has $k = \Theta(\log d)$ levels and a colorful choice has at most d colors. Thus, for $d^2 k + 1 = \Theta(d^2 \log d)$ colors, the pigeonhole principle implies that there is a cell with $d+1$ sets.

Let us consider the running time. One combination step takes $O(d^5)$ time. To compute a set in level i, we have to compute $d+1$ sets in level $i-1$. Hence, computing one set in level $k+1$ takes $d^{O(\log d)}$ time. ☐

Acknowledgements. We would like to extend our thanks to Frédéric Meunier and Pauline Sarrabezolles for interesting discussions on the colorful Carathéodory problem and for hosting us during a research stay at the École Nationale des Ponts et Chaussées.

References

[1] I. Bárány. A generalization of Carathéodory’s theorem. *Discrete Math.*, 40(2–3):141–152, 1982.

[2] I. Bárány and S. Onn. Colourful linear programming and its relatives. *Math. Oper. Res.*, 22(3):550–567, 1997.

[3] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Comput. System Sci.*, 37(1):79–100, 1988.

[4] J. Matoušek. *Lectures on discrete geometry.* Springer, 2002.

[5] F. Meunier and A. Deza. A further generalization of the colourful Carathéodory theorem. In *Discrete geometry and optimization*, volume 69 of *Fields Inst. Commun.*, pages 179–190. Springer, New York, 2013.

[6] F. Meunier and P. Sarrabezolles. Colorful linear programming, Nash equilibrium, and pivots. *arxiv:1409.3436*, 2014.

[7] G. L. Miller and D. R. Sheehy. Approximate centerpoints with proofs. *Comput. Geom.*, 43(8):647–654, 2010.

[8] W. Mulzer and Y. Stein. Complexity of finding nearest colorful polytopes. *Proceedings of the 30th European Workshop on Computational Geometry*, 2014.

[9] C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. System Sci.*, 48(3):498–532, 1994.

[10] K. S. Sarkaria. Tverberg’s theorem via number fields. *Israel J. Math.*, 79(2–3):317–320, 1992.

[11] H. Tverberg. Further generalization of Radon’s theorem. *J. London Math. Soc.*, 43:352–354, 1968.