The Rainbow at the End of the Line — A PPAD Formulation of the Colorful Carathéodory Theorem with Applications

Frédéric Meunier[*] *Wolfgang Mulzer*[†] *Pauline Sarrabezolles*[*] *Yannik Stein*[†]

Abstract

Let $C_1, ..., C_{d+1}$ be $d+1$ point sets in $\mathbb{R}^d$, each containing the origin in its convex hull. A subset C of $\bigcup_{i=1}^{d+1} C_i$ is called a colorful choice (or rainbow) for $C_1, \dots, C_{d+1}$, if it contains exactly one point from each set C_i. The colorful Carathéodory theorem states that there always exists a colorful choice for $C_1, \dots, C_{d+1}$ that has the origin in its convex hull. This theorem is very general and can be used to prove several other existence theorems in high-dimensional discrete geometry, such as the centerpoint theorem or Tverberg's theorem. The colorful Carathéodory problem (ColorfulCarathéodory) is the computational problem of finding such a colorful choice. Despite several efforts in the past, the computational complexity of ColorfulCarathéodory in arbitrary dimension is still open.

We show that ColorfulCarathéodory lies in the intersection of the complexity classes PPAD and PLS. This makes it one of the few geometric problems in PPAD and PLS that are not known to be solvable in polynomial time. Moreover, it implies that the problem of computing centerpoints, computing Tverberg partitions, and computing points with large simplicial depth is contained in PPAD $\cap$ PLS. This is the first nontrivial upper bound on the complexity of these problems.

Finally, we show that our PPAD formulation leads to a polynomial-time algorithm for a special case of ColorfulCarathéodory in which we have only two color classes C_1 and C_2 in d dimensions, each with the origin in its convex hull, and we would like to find a set with half the points from each color class that contains the origin in its convex hull.

1 Introduction

Let $P \subset \mathbb{R}^d$ be a d-dimensional point set. We say P *embraces* a point $\boldsymbol{p} \in \mathbb{R}^d$ or P is $\boldsymbol{p}$-*embracing* if $\boldsymbol{p} \in \operatorname{conv}(P)$, and we say P *ray-embraces* $\boldsymbol{p}$ if $\boldsymbol{p} \in \operatorname{pos}(C)$, where $\operatorname{pos}(P) = \{ \sum_{\boldsymbol{p}\in P} \alpha_{\boldsymbol{p}} \boldsymbol{p} \mid \alpha_{\boldsymbol{p}} \geq 0 \text{ for all } \boldsymbol{p} \in P\}$. Carathéodory's theorem [16, Theorem 1.2.3] states that if P embraces the origin, then there exists a subset $P' \subseteq P$ of size $d+1$ that also embraces the origin. This was generalized by Bárány [4] to the *colorful* setting: let $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ be point sets that each embrace the origin. We call a set $C = \{\boldsymbol{c}_1, \dots, \boldsymbol{c}_{d+1}\}$ a *colorful choice* (or *rainbow*) for $C_1, \dots, C_{d+1}$, if $\boldsymbol{c}_i \in C_i$, for $i = 1, \dots, d+1$. The *colorful* Carathéodory theorem states that there always exists a $\mathbf{0}$-embracing colorful choice that contains the origin in its convex hull. Bárány also gave the following generalization.

[*]Université Paris Est, CERMICS (ENPC), `{frederic.meunier,pauline.sarrabezolles}@enpc.fr`.

[†]Institut für Informatik, Freie Universität Berlin, `{mulzer,yannikstein}@inf.fu-berlin.de`. WM was supported in part by DFG Grants MU 3501/1 and MU 3501/2. YS was supported by the Deutsche Forschungsgemeinschaft within the research training group 'Methods for Discrete Structures' (GRK 1408) and by GIF grant 1161.

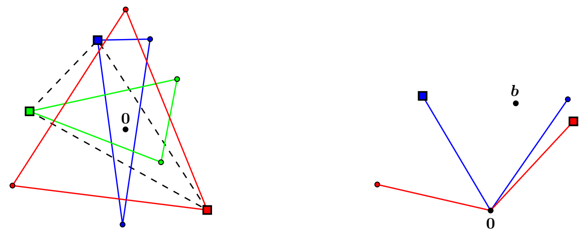

(a) (b)

Fig. 1: (a) Example of the convex version of Theorem 1.1 in two dimensions. (b) Example of the cone version of Theorem 1.1 in two dimensions.

Theorem 1.1 (Colorful Carathéodory Theorem, Cone Version [4])**.** *Let $C_1, \dots, C_d \subset \mathbb{R}^d$ be point sets and $\boldsymbol{b} \in \mathbb{R}^d$ a point with $\boldsymbol{b} \in \operatorname{pos}(C_i)$, for $i = 1, \dots, d$. Then, there is a colorful choice C for $C_1, \dots, C_d$ that ray-embraces $\boldsymbol{b}$.* □

The classic (convex) version of the colorful Carathéodory theorem follows easily from Theorem 1.1: lift the sets $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ to $\mathbb{R}^{d+1}$ by appending a 1 to each element, and set $\boldsymbol{b} = (0, \dots, 0, 1)^T$. See Figure 1 for an example of both versions in two dimensions.

Even though the cone version of the colorful Carathéodory theorem guarantees the existence of a colorful choice that ray-embraces the point $\boldsymbol{b}$, it is far from clear how to find it efficiently. We call this computational problem the *colorful Carathéodory problem* (ColorfulCarathéodory). To this day, settling the complexity of ColorfulCarathéodory remains an intriguing open problem, with a potentially wide range of consequences. We can use linear programming to check in polynomial time whether a given colorful choice ray-embraces a point, so ColorfulCarathéodory lies in *total function NP* (TFNP) [23], the complexity class of total search problems that can be solved in non-deterministic polynomial time. This implies that ColorfulCarathéodory cannot be NP-hard unless NP = coNP [12]. However, the complexity landscape inside TFNP is far from understood, and there exists a rich body of work that studies subclasses of TFNP meant to capture different aspects of mathematical existence proofs, such as the pigeonhole principle (PPP), potential function arguments (PLS, CLS), or various parity arguments (PPAD, PPA, PPADS) [10, 12, 23].

While the complexity of ColorfulCarathéodory remains elusive, related problems are known to be complete for PPAD or for PLS. For example, given $d+1$ point sets $C_1, \dots, C_{d+1} \subset \mathbb{Q}^d$ consisting of two points each and a colorful choice C for $C_1, \dots, C_{d+1}$ that embraces the origin, it is PPAD-complete to find another colorful choice that embraces the origin [20]. Furthermore, given $d+1$ point sets $C_1, \dots, C_{d+1} \subset \mathbb{Q}^d$, we call a colorful choice C for $C_1, \dots, C_{d+1}$ *locally optimal* if the L_1-distance of $\operatorname{conv}(C)$ to the origin cannot be decreased by swapping a point of color i in C with another point from the same color. Then, computing a locally optimal colorful choice is PLS-complete [22].

Understanding the complexity of ColorfulCarathéodory becomes even more interesting in the light of the fact that the colorful Carathéodory theorem plays a crucial role in proving several other prominent theorems in convex geometry, such as Tverberg's theorem [25] (and hence the centerpoint theorem [24]) and the first selection lemma [16, 4]. In fact, these proofs can be interpreted as polynomial time reductions from the respective computational problems, Tverberg, Centerpoint, and SimplicialCenter, to ColorfulCarathéodory. See Section A for more details.

Several approximation algorithms have been proposed for ColorfulCarathéodory. Bárány and Onn [5] describe an exact algorithm that can be stopped early to find a colorful choice whose convex hull is "close" to the origin. More precisely, let $\varepsilon, \rho > 0$ be parameters. We call a set *ε-close* if its convex hull has L_2-distance at most ε to the origin. Given sets $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ such that (i) each C_i contains a ball of radius ρ centered at the origin in its convex hull; and (ii) all points $\boldsymbol{p} \in \bigcup_{i=1}^{d+1} C_i$ fulfill $1 \le \|\boldsymbol{p}\| \le 2$ and can be encoded using L bits, one can find an ε-close colorful choice in time $O(\text{poly}(L, \log(1/\varepsilon), 1/\rho))$ on the Word-Ram with logarithmic costs. For $\varepsilon = 0$, the algorithm actually finds a solution to ColorfulCarathéodory in finite time, and, more interestingly, if $1/\rho = O(\text{poly}(L))$, the algorithm finds a solution to ColorfulCarathéodory in polynomial time. In the same spirit, Barman [6] showed that if the points have constant norm, an ε-close colorful choice can be found by solving $d^{O(1/\varepsilon^2)}$ convex programs. Mulzer and Stein [22] considered a different notion of approximation: a set is called *m-colorful* if it contains at most m points from each C_i. They showed that for all fixed $\varepsilon > 0$, an $\lceil \varepsilon d \rceil$-colorful choice that contains the origin in its convex hull can be found in polynomial time.

Our Results. We provide a new upper bound on the complexity of ColorfulCarathéodory by showing that the problem is contained in PPAD ∩ PLS, implying the first nontrivial upper bound on the computational complexity of computing centerpoints or finding Tverberg partitions.

The traditional proofs of the colorful Carathéodory theorem all proceed through a potential function argument. Thus, it may not be surprising that ColorfulCarathéodory lies in PLS, even though a detailed proof that can deal with degenerate instances requires some care (see Section C). On the other hand, showing that ColorfulCarathéodory lies in PPAD calls for a completely new approach. Even though there are proofs of the colorful Carathéodory theorem that use topological methods usually associated with PPAD (such as certain variants of Sperner's lemma) [11, 13], these proofs involve existential arguments that have no clear algorithmic interpretation. Thus, we present a new proof of the colorful Carathéodory theorem that proceeds similarly as the usual proof for Sperner's lemma [8]. This new proof has an algorithmic interpretation that leads to a formulation of ColorfulCarathéodory as a PPAD-problem.

Finally, we consider the special case of ColorfulCarathéodory that we are given two color classes $C_1, C_2 \subset \mathbb{R}^d$ of d points each and a vector $\boldsymbol{b} \in \mathbb{R}^d$ such that both C_1 and C_2 ray-embrace $\boldsymbol{b}$. We describe an algorithm that solves the following problem in polynomial time: given $k \in [d]$, find a set $C \subseteq C_1 \cup C_2$ with $|C \cap C_1| = k$ and $|C \cap C_2| = d - k$ such that C ray-embraces $\boldsymbol{b}$. Note that this is a special case of ColorfulCarathéodory since we can just take k copies of C_1 and $d - k$ copies of C_2 in a problem instance for ColorfulCarathéodory.

2 Preliminaries

The Complexity Class PPAD. The complexity class *polynomial parity argument in a directed graph* (PPAD) [23] is a subclass of TFNP that contains search problems that can be modeled as follows: let $G = (V, E)$ be a directed graph in which each node has indegree and outdegree at most one. That is, G consists of paths and cycles. We call a node $v \in V$ a *source* if v has indegree 0 and we call v a *sink* if it has outdegree 0. Given a source in G, we want to find another source or sink. By a parity argument, there is an even number of sources and sinks in G and hence another source or sink must exist. However, finding this sink or source is nontrivial since G is defined implicitly and the total number of nodes may be exponential.

More formally, a problem in PPAD is a relation $\mathcal{R}$ between a set $\mathcal{I} \subseteq \{0,1\}^\star$ of *problem instances* and a set $\mathcal{S} \subset \{0,1\}^\star$ of *candidate solutions*. Assume further the following.

- The set $\mathcal{I}$ is polynomial-time verifiable. Furthermore, there is an algorithm that on input $I \in \mathcal{I}$ and $s \in \mathcal{S}$ decides in time $\text{poly}(|I|)$ whether s is a *valid* candidate solution for I. We denote with $\mathcal{S}_I \subseteq \mathcal{S}$ the set of all valid candidate solutions for a fixed instance I.

- There exist two polynomial-time computable functions pred and succ that define the edge set of G as follows: on input $I \in \mathcal{I}$ and $s \in \mathcal{S}_I$, pred and succ return a valid candidate solution from $\mathcal{S}_I$ or $\perp$. Here, $\perp$ means that v has no predecessor/successor.

- There is a polynomial-time algorithm that returns for each instance I a valid candidate solution $s \in \mathcal{S}_I$ with $\mathrm{pred}(s) = \perp$. We call s the *standard source*.

Now, each instance $I \in \mathcal{I}$ defines a graph $G_I = (V, E)$ as follows. The set of nodes V is the set of all valid candidate solutions $\mathcal{S}_I$ and there is a directed edge from u to v if and only if $v = \mathrm{succ}(u)$ and $u = \mathrm{pred}(v)$. Clearly, each node in G_I has indegree and outdegree at most one. The relation $\mathcal{R}$ consists of all tuples (I, s) such that s is a sink or source other than the standard source in G_I.

The definition of a PPAD-problem suggests a simple algorithm, called the *standard algorithm*: start at the standard source and follow the path until a sink is reached. This algorithm always finds a solution but the length of the traversed path may be exponential in the size of the input instance.

Polyhedral Complexes and Subdivisions. We call a finite set of polyhedra $\mathcal{P}$ in $\mathbb{R}^d$ a *polyhedral complex* if and only if (i) for all polyhedra $f \in \mathcal{P}$, all faces of f are contained in $\mathcal{P}$; and (ii) for all $f, f' \in \mathcal{P}$, the intersection $f \cap f'$ is a face of both. Note that the first requirement implies that $\emptyset \in \mathcal{P}$. Furthermore, we say $\mathcal{P}$ has *dimension* k if there exists some polyhedron $f \in \mathcal{P}$ with $\dim f = k$ and all other polyhedra in $\mathcal{P}$ have dimension at most k. We call $\mathcal{P}$ a *polytopal complex* if it is a polyhedral complex and all elements are polytopes. Similarly, we say $\mathcal{P}$ is a *simplicial complex* if it is a polytopal complex whose elements are simplices. Finally, we say $\mathcal{P}$ *subdivides* a set $Q \subseteq \mathbb{R}^d$ if $\bigcup_{f \in \mathcal{P}} f = Q$. For more details, see [27, Section 5.1].

Linear Programming. Let $A \in \mathbb{R}^{d \times n}$ be a matrix and F a set of column vectors from A. Then, we denote with $\mathrm{ind}\,(F) \subseteq [n]$ the set of column indices in F and for an index set $I \subseteq [n]$, we denote with A_I the submatrix of A that consists of the columns indexed by I. Similarly, for a vector $\boldsymbol{c} \in \mathbb{R}^n$ and an index set $I \subset [n]$, we denote with $\boldsymbol{c}_I$ the subvector of $\boldsymbol{c}$ with the coordinates indexed by I. Now, let L' denote a system of linear equations

$$L' : A\boldsymbol{x} = \boldsymbol{b},$$

where $A \in \mathbb{Q}^{d \times n}$, $\boldsymbol{b} \in \mathbb{Q}^d$ and $\mathrm{rank}(A) = k$. By multiplying with the least common denominator, we may assume in the following that $A \in \mathbb{Z}^{d \times n}$ and $\boldsymbol{b} \in \mathbb{Z}^d$. We call a set of k linearly independent column vectors B of A a *basis* and we say that A is *non-degenerate* if $k = d$ and for all bases B of A, no coordinate of the corresponding solution $\boldsymbol{x}_{\mathrm{ind}(B)}$ is 0. In particular, if L' is non-degenerate, then $\boldsymbol{b}$ is not contained in the linear span of any set of $d' < d$ column vectors from A and hence if $d > n$, the linear system L' has no solution. In the following, we assume that L' is non-degenerate and that $d \leq n$.

We denote with L the linear program obtained by extending the linear system L' with the constraints $\boldsymbol{x} \geq \boldsymbol{0}$ and with a cost vector $\boldsymbol{c} \in \mathbb{Q}^n$:

$$L : \min \boldsymbol{c}^T \boldsymbol{x} \text{ subject to } A\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{x} \geq \boldsymbol{0}.$$

We say a set of column vectors B is a *basis* for L if B is a basis for L'. Let $\boldsymbol{x} \in \mathbb{R}^n$ be the corresponding solution, i.e., let $\boldsymbol{x}$ be such that $A\boldsymbol{x} = \boldsymbol{b}$ and $x_i = 0$ for $i \in [n] \setminus \mathrm{ind}\,(B)$. We call $\boldsymbol{x}$ a *basic feasible solution*, and B a *feasible basis*, if $\boldsymbol{x} \geq \boldsymbol{0}$. Furthermore, we say L is *non-degenerate* if for all feasible bases B, the corresponding basic feasible solutions have strictly positive values in the coordinates of B. Now, let $R = [n] \setminus \text{ind}(B)$ be the column indices not in B. The *reduced cost vector* $\boldsymbol{r}_{B,\boldsymbol{c}} \in \mathbb{Q}^{n-d}$ with respect to B and $\boldsymbol{c}$ is then defined as

$$\boldsymbol{r}_{B,\boldsymbol{c}} = \boldsymbol{c}_R - \left(A^{-1}_{\text{ind}(B)} A_R\right)^T \boldsymbol{c}_{\text{ind}(B)}. \tag{1}$$

It is well-known that B is optimal for $\boldsymbol{c}$ if and only if $\boldsymbol{r}_{B,\boldsymbol{c}}$ is non-negative in all coordinates [18]. For technical reasons, we consider in the following the *extended reduced cost vector* $\hat{\boldsymbol{r}}_{B,\boldsymbol{c}} \in \mathbb{Q}^n$ that has a 0 in dimensions $\text{ind}(B)$ and otherwise equals $\boldsymbol{r}_{B,\boldsymbol{c}}$ to align the coordinates of the reduced cost vector with the column indices in A. More formally, we set

$$(\hat{\boldsymbol{r}}_{B,\boldsymbol{c}})_j = \begin{cases} 0 & \text{if } j \in \text{ind}(B), \text{ and} \\ (\boldsymbol{r}_{B,\boldsymbol{c}})_{j'} & \text{otherwise,} \end{cases}$$

where j' is the rank of j in R, that is, $(\boldsymbol{r}_{B,\boldsymbol{c}})_{j'}$ is the coordinate of $\boldsymbol{r}_{B,\boldsymbol{c}}$ that corresponds to the j'th non-basis column with column index j in A.

Geometrically, the feasible solutions for the linear program L define an $(n-d)$-dimensional polyhedron $\mathcal{P}$ in $\mathbb{R}^n$. Since L is non-degenerate, $\mathcal{P}$ is simple. Let $f \subseteq \mathcal{P}$ be a k-face of $\mathcal{P}$. Then, f has an associated set $\text{supp}(f) \subseteq [n]$ of k column indices such that f consists precisely of the feasible solutions for the linear program $A_{\text{supp}(f)}\boldsymbol{x'} = \boldsymbol{b}$, $\boldsymbol{x'} \geq \boldsymbol{0}$, lifted to $\mathbb{R}^n$ by setting the coordinates with indices not in $\text{supp}(f)$ to 0. We call $\text{supp}(f)$ the *support* of f and we say the columns in $A_{\text{supp}(f)}$ *define* f. Furthermore, for all subfaces $\check{f} \subseteq f$, we have $\text{supp}\left(\check{f}\right) \subseteq \text{supp}(f)$ and in particular, all bases that define vertices of f are d-subsets of columns from $A_{\text{supp}(f)}$.

Moreover, we say a nonempty face $f \subseteq \mathcal{P}$ is *optimal* for a cost vector $\boldsymbol{c}$ if all points in f are optimal for $\boldsymbol{c}$. We can express this condition using the reduced cost vector. Let B be a basis for a vertex in f. Then f is optimal for $\boldsymbol{c}$ if and only if

$$(\hat{\boldsymbol{r}}_{B,\boldsymbol{c}})_j = 0 \text{ for } j \in \text{supp}(f), \text{ and } (\hat{\boldsymbol{r}}_{B,\boldsymbol{c}})_j \leq 0 \text{ otherwise.}$$

3 Overview of the PPAD-Formulation

We give a new constructive proof of the cone version of the colorful Carathéodory theorem based on Sperner's lemma. Using this, we can obtain a PPAD-formulation of ColorfulCarathéodory, by adapting Papadimitriou's formulation of Sperner's lemma as a PPAD problem.

Recall the statement of Sperner's lemma: let $\mathcal{S}$ be a simplicial subdivision of the d-dimensional standard simplex $\Delta^d = \text{conv}(\boldsymbol{e}_1, \dots, \boldsymbol{e}_{d+1}) \subset \mathbb{R}^{d+1}$, where $\boldsymbol{e}_i$ is the ith canonical basis vector. We call a function λ that assigns to each vertex in $\mathcal{S}$ a label from $[d+1]$ a *Sperner labeling* if for each vertex $\boldsymbol{v}$ of $\mathcal{S}$ contained in $\text{conv}(\boldsymbol{e}_{i_1}, \dots, \boldsymbol{e}_{i_k})$, we have $\lambda(\boldsymbol{v}) \in \{i_1, \dots, i_k\}$, for all $\{i_1, \dots, i_k\} \subseteq [d+1], k \in [d+1]$. For a simplex $\sigma \in \mathcal{S}$, we set $\lambda(\sigma)$ to be the set of labels of the vertices of σ. We call σ *fully-labeled* if $\lambda(\sigma) = [d+1]$.

Theorem 3.1 (Strong Sperner's Lemma [8])**.** *The number of fully-labeled simplices is odd.*

Now suppose we are given an instance $I = (C_1, \dots, C_d, \boldsymbol{b})$ of (the cone version of) Colorful-Carathéodory, where $\boldsymbol{b} \in \mathbb{R}^d$, $\boldsymbol{b} \neq \boldsymbol{0}$, and each $C_i \subset \mathbb{Q}^d$, $i \in [d]$, ray-embraces $\boldsymbol{b}$. In Section B, we show that we can assume w.l.o.g. that each set C_i has size d. We now describe how to define a simplicial complex $\mathcal{S}$ and a Sperner labeling λ for I such that a fully labeled simplex will encode a colorful choice that contains the vector $\boldsymbol{b}$ in its positive span.

In the following, we call $\mathbb{R}^d$ the *parameter space* and a vector $\boldsymbol{\mu} \in \mathbb{R}^d$ a *parameter vector*. We define a family of linear programs $\{L^{\text{CC}}_{\boldsymbol{\mu}} \mid \boldsymbol{\mu} \in \mathbb{R}^d\}$, where each linear program $L^{\text{CC}}_{\boldsymbol{\mu}}$ has the same

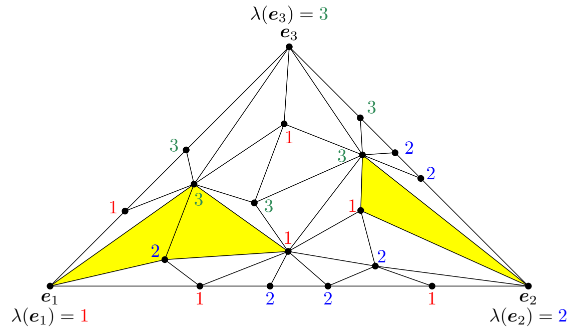

Fig. 2: An example of Sperner's lemma in two dimensions. The fully-labeled simplices are marked yellow.

constraints and differs only in its cost vector $\boldsymbol{c}_{\boldsymbol{\mu}}$. The cost vector $\boldsymbol{c}_{\boldsymbol{\mu}}$ is defined by a linear function in $\boldsymbol{\mu} \in \mathbb{R}^d$. Let $A = (C_1\ C_2\ \ldots\ C_d) \in \mathbb{Q}^{d \times d^2}$ be the matrix that has the vectors from C_1 in the first d columns, the vectors from C_2 in the second d columns, and so on. Then, we denote with $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$ the linear program

$$L^{\mathrm{CC}}_{\boldsymbol{\mu}} : \min \boldsymbol{c}^T_{\boldsymbol{\mu}} \boldsymbol{x}, \text{ subject to } A\boldsymbol{x} = \boldsymbol{b}, \boldsymbol{x} \geq \boldsymbol{0}, \tag{2}$$

and we denote with $\mathcal{P}^{\mathrm{CC}} \subset \mathbb{R}^{d^2}$ the polyhedron that is defined by the linear system L^{CC}. We can think of the ith coordinate of the parameter vector $\boldsymbol{\mu} \in \mathbb{R}^d$ as the weight of color i, i.e., the costs of columns from A with color i decrease if $(\boldsymbol{\mu})_i$ increases. To each face f of P, we assign the set of parameter vectors $\Phi(f) \subset \mathbb{R}^d$ such that for all $\boldsymbol{\mu} \in \Phi(f)$, the face f is optimal for the linear program $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$ that has L^{CC} as constraints and $\boldsymbol{c}_{\boldsymbol{\mu}}$ as cost vector. We call $\Phi(f)$ the *parameter region* of f. The cost vector is designed to control the colors that appear in the support of optimal faces for a specific subset of parameter vectors. Let $\mathcal{M} = \left\{ \boldsymbol{\mu} \in \mathbb{R}^d \,\middle|\, \boldsymbol{\mu} \geq \boldsymbol{0},\ \|\boldsymbol{\mu}\|_\infty = 1 \right\}$ denote the faces of the unit cube in which at least one coordinate is set to 1. Then, no face f that is assigned to a parameter vector $\boldsymbol{\mu} \in \mathcal{M}$ with $(\boldsymbol{\mu})_{i^\times} = 0$ has a column from A with color $i^\times$ in its defining set $A_{\mathrm{supp}(f)}$. This property will become crucial when we define a Sperner labeling later on. Now, we define a polyhedral subcomplex $\mathcal{F}$ of $\mathcal{P}^{\mathrm{CC}}$ that consists of all faces f of $\mathcal{P}^{\mathrm{CC}}$ such that $\Phi(f) \cap \mathcal{M} \neq \emptyset$. Furthermore, the intersections of the parameter regions with $\mathcal{M}$ induce a polytopal complex $\mathcal{Q}$ that is in a dual relationship to $\mathcal{F}$. By performing a central projection with the origin as center of $\mathcal{Q}$ onto the standard simplex Δ^{d-1}, we obtain a polytopal subdivision $\mathcal{Q}_\Delta$ of Δ^{d-1}. To get the desired *simplicial* subdivision of Δ^{d-1}, we take the barycentric subdivision $\operatorname{sd} \mathcal{Q}_\Delta$ of $\mathcal{Q}_\Delta$.

We construct a Sperner labeling λ for $\operatorname{sd} \mathcal{Q}_\Delta$ as follows: let $\boldsymbol{v}$ be a vertex in $\operatorname{sd} \mathcal{Q}_\Delta$, and let f be the face of $\mathcal{F}$ that corresponds to $\boldsymbol{v}$. Then, we set $\lambda(\boldsymbol{v}) = i$ if the ith color appears most often in the support of f. The color controlling property of the cost function $c_{\boldsymbol{\mu}}$ then implies that λ is a Sperner labeling. Furthermore, using the properties of the barycentric subdivision and the correspondence between $\mathcal{Q}_\Delta$ and $\mathcal{F}$, we can show that one vertex of a fully-labeled $(d-1)$-simplex in $\operatorname{sd} \mathcal{Q}_\Delta$ encodes a colorful feasible basis of the ColorfulCarathéodory instance I. This concludes a new constructive proof of the colorful Carathéodory theorem using Sperner's lemma.

To show that ColorfulCarathéodory is in PPAD however, we need to be able to traverse $\operatorname{sd} \mathcal{Q}_\Delta$ efficiently. For this, we introduce a combinatorial encoding of the simplices in $\mathcal{Q}_\Delta$ that represents neighboring simplices in a similar manner. Furthermore, we describe how to generalize the orientation used in the PPAD formulation of 2D-Sperner [23] to our setting. This finally shows that ColorfulCarathéodory is in PPAD.

To ensure that the complexes that appear in our algorithms are sufficiently generic, we prove several perturbation lemmas that give a deterministic way of achieving this. Our PPAD-formulation also shows that the special case of ColorfulCarathéodory involving two colors can be solved in polynomial time. Indeed, we will see that in this case the polytopal complex $\mathcal{Q}_\Delta$ can be made 1-dimensional. Then, binary search can be used to find a fully-labeled simplex in $\mathcal{Q}_\Delta$. In order to prove that the binary search terminates after a polynomial number of steps, we use methods similar to our perturbation techniques to obtain a bound on the length of the 1-dimensional fully-labeled simplex.

4 The Colorful Carathéodory Problem is in PPAD

As before, let $I = (C_1, \dots, C_d, \boldsymbol{b})$ denote an instance for the cone version of ColorfulCarathéodory. Our formulation of ColorfulCarathéodory as a PPAD-problem requires I to be in general position. In particular, we assume that (P1) all color classes $C_i \subset \mathbb{Z}^d$ consist of d points and all points have integer coordinates. Furthermore, we assume that (P2) there exist no subset $P \subset \bigcup_{i=1}^d C_i$ of size $d-1$ that ray-embraces $\boldsymbol{b}$. We show in Section B how to ensure the properties by an explicit deterministic perturbation of polynomial bit-complexity.

4.1 The Polytopal Complex

Let $N = d!m^d$, where m is the largest absolute value that appears in A and $\boldsymbol{b}$ (see Lemma D.1). Then, we define $\boldsymbol{c_\mu} \in \mathbb{R}^{d^2}$ as

$$(\boldsymbol{c_\mu})_j = 1 + (1 - (\boldsymbol{\mu})_i)\, dN^2 + \varepsilon^j, \tag{3}$$

where $j \in [d^2]$, i is the color of the jth column in A, and $0 < \varepsilon \leq N^{-3}$ is a suitable perturbation that ensures non-degeneracy of the reduced costs (see [7]). As stated in the overview, the cost function controls the colors in the support of the optimal faces for parameter vectors in $\mathcal{M}$. The proof of the following lemma can be found in Section D.

Lemma 4.1. *Let $i^\times \in [d]$ be a color and let $\boldsymbol{\mu} \in \mathcal{M}$ be a parameter vector with $\boldsymbol{\mu}_{i^\times} = 0$. Furthermore, let $B^\star$ be an optimal feasible basis for $L_{\boldsymbol{\mu}}^{\mathrm{CC}}$. Then, $B^\star \cap C_{i^\times} = \emptyset$.*

We denote for a face $f \subseteq \mathcal{P}^{\mathrm{CC}}$, $f \neq \emptyset$, with $\Phi(f) = \{\boldsymbol{\mu} \in \mathbb{R}^d \mid f \text{ is optimal for } L_{\boldsymbol{\mu}}\}$ the set of all parameter vectors for which f is optimal. We call this the *parameter region* for f. Using the reduced cost vector, we can express $\Phi(f)$ as solution space to the following linear system, where B is a feasible basis of some vertex of f and the d coordinates of the parameter vector $\boldsymbol{\mu}$ are the variables:

$$L_{B,f}^{\Phi} : (\hat{\boldsymbol{r}}_{B,\boldsymbol{c_\mu}})_j = 0 \text{ for } j \in \operatorname{supp}(f) \setminus \operatorname{ind}(B) \text{ and } (\hat{\boldsymbol{r}}_{B,\boldsymbol{c_\mu}})_j \leq 0 \text{ for } [d^2] \setminus \operatorname{supp}(f). \tag{4}$$

Then, we define $\mathcal{F}$ as the set of all faces that are optimal for some parameter vector in $\mathcal{M}$:

$$\mathcal{F} = \left\{ f \,\middle|\, f \text{ is a face of } \mathcal{P}^{\mathrm{CC}},\ \Phi(f) \cap \mathcal{M} \neq \emptyset \right\}.$$

By definition, $\mathcal{F} \cup \{\emptyset\}$ is a polyhedral subcomplex of $\mathcal{P}^{\mathrm{CC}}$. The intersections of the parameter regions with faces of $\mathcal{M}$ induce a subdivision $\mathcal{Q}$ of $\mathcal{M}$:

$$\mathcal{Q} = \{\Phi(f) \cap g \mid f \in \mathcal{F},\ g \text{ is a face of } \mathcal{M}\}.$$

In Section D, we show that $\mathcal{Q}$ is a $(d-1)$-dimensional polytopal complex. Next, we construct $\mathcal{Q}_\Delta$ through a central projection with the origin as center of $\mathcal{Q}$ onto the $(d-1)$-dimensional standard simplex $\Delta \subset \mathbb{R}^d$. It is easy to see that this projection is a bijection. For a parameter vector $\boldsymbol{\mu} \in \mathbb{R}^d$, we denote with $\Delta(\boldsymbol{\mu}) = \boldsymbol{\mu}/\|\boldsymbol{\mu}\|_1$ its projection onto Δ. Similarly, we denote with $\mathcal{M}(\boldsymbol{\mu}) = \boldsymbol{\mu}/\|\boldsymbol{\mu}\|_\infty$ the projection of $\boldsymbol{\mu}$ onto $\mathcal{M}$ and we use the same notation to denote the element-wise projection of sets. Then, we can write the projection $\mathcal{Q}_\Delta$ of $\mathcal{Q}$ onto Δ as $\mathcal{Q}_\Delta = \{\Delta(q) \mid q \in \mathcal{Q}\}$. Furthermore, let $\mathcal{S} = \{\Delta(g) \mid g \text{ is a face of } \mathcal{M}\}$ denote the projections of the faces of $\mathcal{M}$ onto Δ. For $f \in \mathcal{F}$, let $\Phi_\Delta(f) = \Delta(\Phi(f) \cap \mathcal{M})$ denote the projection of all parameter vectors in $\mathcal{M}$ for which f is optimal onto Δ. Please refer to Table 1 on Page 14 for an overview of the current and future notation. The following results are proved in Section D.

Lemma 4.2. *Let $q \neq \emptyset$ be an element from $\mathcal{Q}_\Delta$. Then, there exists unique pair (f, g) where f is a face of $\mathcal{F}$ and g is a face of $\mathcal{S}$ such that $q = \Phi_\Delta(f) \cap g$. Moreover, q is a simple polytope of dimension $\dim g - \dim f$ and, if $\dim q > 0$, the set of facets of q can be written as*

$$\left\{\Phi_\Delta(f) \cap \check{g} \neq \emptyset \,\middle|\, \check{g} \text{ is a facet of } g\right\} \cup \left\{\Phi_\Delta\left(\hat{f}\right) \cap g \neq \emptyset \,\middle|\, f \text{ is a facet of } \hat{f} \in \mathcal{F}\right\}.$$

Lemma 4.3. *The set $\mathcal{Q}_\Delta$ is a $(d-1)$-dimensional polytopal complex that decomposes Δ.* ◻

4.2 The Barycentric Subdivision

The *barycentric subdivision* [17, Definition 1.7.2] is a well-known method to subdivide a polytopal complex into simplices. We define $\operatorname{sd} \mathcal{Q}_\Delta$ as the set of all simplices $\operatorname{conv}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_k)$, $k \in [d]$, such that there exists a chain $q_0 \subset \cdots \subset q_k$ of polytopes in $\mathcal{Q}_\Delta$ with $\dim q_{i-1} < \dim q_i$ and such that $\boldsymbol{v}_i$ is the barycenter of q_i for $i \in [k]$. We define the label of a vertex $\boldsymbol{v} \in \operatorname{sd} \mathcal{Q}_\Delta$ as follows. By Lemma 4.2, there exists a unique pair $f \in \mathcal{F}$ and $g \in \mathcal{S}$ with $\boldsymbol{v} = \Phi_\Delta(f) \cap g$. Then, the label $\lambda(\boldsymbol{v})$ of $\boldsymbol{v}$ is defined as

$$\lambda(\boldsymbol{v}) = \arg\max_{i \in [d]} \left|\operatorname{ind}(C)_i \cap \operatorname{supp}(f)\right|. \tag{5}$$

In case of a tie, we take the smallest $i \in [d]$ that achieves the maximum. Lemma 4.1 implies that $\lambda(\cdot)$ is a Sperner labeling of $\operatorname{sd} \mathcal{Q}_\Delta$. In fact, λ is a Sperner labeling for any fixed simplicial subdivision of Δ. Now, Theorem 3.1 guarantees the existence of a $(d-1)$-simplex $\sigma \in \operatorname{sd} \mathcal{Q}_\Delta$ whose vertices have all d possible labels. The next lemma shows that then one of the vertices of σ defines a solution to the ColorfulCarathéodory instance. Here, we use specific properties of the barycentric subdivision.

Lemma 4.4. *Let $\sigma \in \operatorname{sd} \mathcal{Q}_\Delta$ be a fully-labeled $(d-1)$-simplex and let $\boldsymbol{v}_{d-1}$ denote the vertex of σ that is the barycenter of a $(d-1)$-face $q_{d-1} = \Phi_\Delta(f_{d-1}) \cap g_{d-1} \in \mathcal{Q}_\Delta$, where $f_{d-1} \in \mathcal{F}$ and $g_{d-1} \in \mathcal{S}$. Then, the columns from $A_{\operatorname{supp}(f_{d-1})}$ are a colorful choice that ray-embraces $\boldsymbol{b}$.*

Our discussion up to now already yields a new Sperner-based proof of the colorful Carathéodory theorem. However, in order to show that ColorfulCarathéodory $\in$ PPAD, we need to replace the invocation of Theorem 3.1 by a PPAD-problem. Note that it is not possible to use the formulation of Sperner from [23, Theorem 2] directly, since it is defined for a fixed simplicial subdivision of the standard simplex. In our case, the simplicial subdivision of Δ depends on the input instance. In the following, we generalize the PPAD formulation of Sperner in [23] to $\mathcal{Q}_\Delta$ by mimicking the proof of Theorem 3.1. For this, we need to be able to find simplices in $\operatorname{sd} \mathcal{Q}_\Delta$ that share a given facet. We begin with a simple encoding of simplices in $\operatorname{sd} \mathcal{Q}_\Delta$ that allows us to solve this problem completely combinatorially.

We first show how to encode a polytope $q \in \mathcal{Q}_\Delta$. By Lemma 4.2, there exists a unique pair of faces $f \in \mathcal{F}$ and $g \in S$ such that $q = \Phi_\Delta(f) \cap g$. Since $\mathcal{M}(g)$ is a face of the unit cube, the value of

$d - \dim g$ coordinates in $\mathcal{M}(g)$ is fixed to either 0 or 1. Let $I_j \subseteq [d]$, $j = 0, 1$, denote the indices of the coordinates that are fixed to j. Then, the encoding of q is defined as $\mathrm{enc}\,(q) = (\mathrm{supp}\,(f), I_0, I_1)$. We use this to define an encoding of the simplices in $\mathcal{Q}_\Delta$ as follows. Let $\sigma \in \mathcal{Q}_\Delta$ be a k-simplex and let $q_0 \subset \cdots \subset q_k$ be the corresponding face chain in $\mathcal{Q}_\Delta$ such that the ith vertex of σ is the barycenter of q_i. Then, the encoding $\mathrm{enc}\,(\sigma)$ is defined as

$$\mathrm{enc}\,(\sigma) = (\mathrm{enc}\,(q_0), \ldots, \mathrm{enc}\,(q_k))\,. \tag{6}$$

In the proof of Theorem 3.1, we traverse only a subset of simplices in the simplicial subdivision, namely $(k-1)$-simplices that are contained in the face $\Delta_{[k]} = \mathrm{conv}\{\boldsymbol{e}_i \mid i \in [k]\}$ of Δ for $k \in [d]$. Let $\Sigma_k = \left\{\sigma \in \mathrm{sd}\,\mathcal{Q}_\Delta \,\middle|\, \dim(\sigma) = k-1,\ \sigma \subseteq \Delta_{[k]}\right\}$ denote the set of $(k-1)$-simplices in $\mathrm{sd}\,\mathcal{Q}_\Delta$ that are contained in the $(k-1)$-face, where $k \in [d]$, and let $\Sigma = \bigcup_{k=1}^{d} \Sigma_k$ be the collection of all those simplices. In the following, we give a precise characterization of the encodings of the simplices in Σ_k. For two disjoint index sets $I_0, I_1 \subseteq [d]$, we denote with $g(I_0, I_1) = \{\boldsymbol{\mu} \in \mathcal{M} \mid j = 0, 1,\ (\boldsymbol{\mu})_i = j \text{ for } i \in I_j\}$ the face of $\mathcal{M}$ that we obtain by fixing the coordinates in dimensions $I_0 \cup I_1$. Let now $T = (Q_0, \ldots, Q_{k-1})$, $k \in [d-1]$, be a tuple, where $Q_i = \left(S^{(i)}, I_0^{(i)}, I_1^{(i)}\right)$, $S^{(i)} \subset [d^2]$, and $I_0^{(i)}, I_1^{(i)}$ are disjoint subsets of $[d]$ with $I_1^{(i)} \neq \emptyset$ for $i \in [k-1]_0$. We say T is *valid* if and only if T has the following properties.

(i) We have $I_0^{(k-1)} = [d] \setminus [k]$, $\left|I_1^{(k-1)}\right| = 1$, and the columns in $A_{S^{(k-1)}}$ are a feasible basis for a vertex f. Moreover, the intersection $\Phi(f) \cap g\left(I_0^{(k-1)} \cup I_1^{(k-1)}\right)$ is nonempty.

(ii) For all $i \in [k-1]$, we either have

  (ii.a) $I_0^{(i-1)} = I_0^{(i)}$, $I_1^{(i-1)} = I_1^{(i)}$, and $S^{(i-1)} = S^{(i)} \cup \{a_{i-1}\}$ for some index $a_{i-1} \in [d^2] \setminus S^{(i)}$,

  (ii.b) or $S^{(i-1)} = S^{(i)}$ and there is an index $j_{i-1} \in [d] \setminus \left(I_0^{(i)} \cup I_1^{(i)}\right)$ such that either $I_0^{(i-1)} = I_0^{(i)}$ and $I_1^{(i-1)} = I_1^{(i)} \cup \{j_{i-1}\}$, or $I_1^{(i-1)} = I_1^{(i)}$ and $I_0^{(i-1)} = I_0^{(i)} \cup \{j_{i-1}\}$.

Lemma 4.5. *For $k \in [d]$, the function* $\mathrm{enc}\,(\cdot)$ *restricted to the simplices in Σ_k is a bijection from Σ_k to the set of valid k-tuples.*

Using our characterization of encodings as valid tuples, it becomes an easy task to check whether a given candidate encoding corresponds to a simplex in Σ.

Lemma 4.6. *Let $T = (Q_0, \ldots, Q_{k-1})$, $k \in [d-1]$, be a tuple, where $Q_i = \left(S^{(i)}, I_0^{(i)}, I_1^{(i)}\right)$, $S^{(i)} \subset [d^2]$, and $I_0^{(i)}, I_1^{(i)}$ are disjoint subsets of $[d]$ with $I_1^{(i)} \neq \emptyset$ for $i \in [k-1]_0$. Then, we can check in polynomial time whether T is a valid k-tuple.*

In Section E, we show that simplices in Σ that share a facet have similar encodings that differ only in one element of the encoding tuples. Using this fact, we can traverse Σ efficiently by manipulating the respective encodings.

Lemma 4.7. *Let $\sigma \in \Sigma_k$ be a simplex and let $q_0 \subset \cdots \subset q_{k-1}$ be the corresponding face chain in $\mathcal{Q}_\Delta$ such that the ith vertex $\boldsymbol{v}_i$ of σ is the barycenter of q_i, where $k \in [d]$ and $i \in [k-1]_0$. Then, we can solve the following problems in polynomial time: (i) Given* $\mathrm{enc}\,(\sigma)$ *and i, compute the encoding of the simplex $\sigma' \in \Sigma_k$ that shares the facet* $\mathrm{conv}\,\{\boldsymbol{v}_j \mid j \in [k-1]_0,\ j \neq i\}$ *with σ or state that there is none; (ii) Assuming that $k < d$ and given* $\mathrm{enc}\,(\sigma)$*, compute the encoding of the simplex $\hat{\sigma} \in \Sigma_{k+1}$ that has σ as facet; and (iii) Assuming that $k > 1$ and given* $\mathrm{enc}\,(\sigma)$*, compute the encoding of the simplex $\check{\sigma} \in \Sigma_{k-1}$ that is a facet of σ or state that there is none.*

4.3 The PPAD graph

Using our tools from the previous sections, we now describe the PPAD graph $G = (V, E)$ for the ColorfulCarathéodory instance. The definition of G follows mainly the ideas from the formulation of Sperner as a PPAD-problem [23, Theorem 2] and the proof of Theorem 3.1.

The graph has one node per simplex in Σ that has all labels or all but the largest possible label. That is, we have one node for each $(k-1)$-simplex σ in Σ_k with $[k-1] \subseteq \lambda(\sigma)$. Two simplices are connected by an edge if one simplex is the facet of the other or if both simplices share a facet that has all but the largest possible label. More formally, for $k \in [d]$, we set $V_k = \{\text{enc}(\sigma) \mid \sigma \in \Sigma_k, [k-1] \subseteq \lambda(\sigma)\}$, the set of all encodings for $(k-1)$-simplices in Σ_k whose vertices have all or all but the largest possible label. Then, V is the union of all V_k for $k \in [d]$. There are two types of edges: edges within a set V_k, $k \in [d]$, and edges connecting nodes from V_k to nodes in V_{k-1} and V_{k+1}. Let $\text{enc}(\sigma), \text{enc}(\sigma')$ be two vertices in V_k for some $k \in [d]$. Then, there is an edge between $\text{enc}(\sigma)$ and $\text{enc}(\sigma')$ if the encoded simplices $\sigma, \sigma' \in \Sigma_k$ share a facet $\check{\sigma}$ with $\lambda(\check{\sigma}) = [k-1]$, i.e., both simplices are connected by a facet that has all but the largest possible label. Now, let $\text{enc}(\sigma) \in V_k$ and $\text{enc}(\sigma') \in V_{k+1}$ for some $k \in [d-1]$. Then, there is an edge between $\text{enc}(\sigma)$ and $\text{enc}(\sigma')$ if $\lambda(\sigma) = [k]$ and σ is a facet of σ'. In the next lemma, we show that G consists only of paths and cycles. Please see Section F for the proof.

Lemma 4.8. *Let* $\text{enc}(\sigma) \in V$ *be a node. If* $\text{enc}(\sigma) \in V_1$ *or* $\text{enc}(\sigma) \in V_d$ *with* $\lambda(\sigma) = [d]$*, then* $\deg \text{enc}(\sigma) = 1$*. Otherwise,* $\deg \text{enc}(\sigma) = 2$*.*

This already shows that ColorfulCarathéodory $\in$ PPA. By generalizing the orientation from [23] to our setting, we obtain a function dir that orients the edges of G such that only vertices with degree one in G are sinks or sources in the oriented graph. In Section F, we show how to compute this function in polynomial time. This finally yields our main result.

Theorem 4.9. ColorfulCarathéodory*,* Centerpoint*,* Tverberg*, and* SimplicialCenter *are in* PPAD $\cap$ PLS*.*

Proof. We give a formulation of ColorfulCarathéodory as PPAD-problem. See Section C for a formulation of ColorfulCarathéodory as PLS-problem. Using the classic proofs discussed in Section A, this then also implies the statement for the other problems.

The set of problem instances $\mathcal{I}$ consists of all tuples $I = (C_1, \ldots, C_d, \boldsymbol{b})$, where $d \in \mathbb{N}$, the set $C_i \subset \mathbb{Q}^d$ ray-embraces $\boldsymbol{b} \in \mathbb{Q}^d$ and $\boldsymbol{b} \neq \mathbf{0}$. Let $I^{\approx} = (C_1^{\approx}, \ldots, C_d^{\approx}, \boldsymbol{b}^{\approx})$ denote then the ColorfulCarathéodory instance that we obtain by applying our perturbation techniques to I (see Section B). Then, $I^{\approx}$ has the general position properties (P1) and (P2). The set of candidate solutions $\mathcal{S}$ consists of all tuples $(Q_0, \ldots, Q_{k-1})$, where $k \in \mathbb{N}$ and Q_i is a tuple $\left(S^{(i)}, I_0^{(i)}, I_1^{(i)}\right)$ with $S^{(i)}, I_0^{(i)}, I_1^{(i)} \subset \mathbb{N}$. Furthermore, $\mathcal{S}$ contains all d-subsets $C \subset \mathbb{Q}^d$ for $d \in \mathbb{N}$. We define the set of valid candidate solutions $\mathcal{S}_I$ for the instance I to be the set of all valid k-tuples with respect to the instance $I^{\approx}$ and the set of all colorful choices with respect to I that ray-embrace $\boldsymbol{b}$, where $k \in [d]$. Let $s \in \mathcal{S}$ be a candidate solution. If it is a tuple, we first use the algorithm from Lemma 4.6 to check in polynomial time in the length of $I^{\approx}$ and hence in the length of I whether $s \in \mathcal{S}_I$. If affirmative, we check whether the simplex has all or all but the largest possible label. Using the encoding, this can be carried out in polynomial time. If s is a set of points, we can determine in polynomial time with linear programming whether the points in s ray-embrace $\boldsymbol{b}$.

We set as standard source the 0-simplex $\{\boldsymbol{e}_1\}$. We can assume without loss of generality that $\{\boldsymbol{e}_1\}$ is a source (otherwise we invert the orientation).

Given a valid candidate solution $s \in \mathcal{S}_I$, we compute its predecessor and successor with the algorithms from Lemma 4.7 and the orientation function discussed above, with one modification:

if a node $s \in V$ is a source different from the standard source in the graph G, it encodes by the above discussion a colorful choice $C^{\approx}$ that ray-embraces $\boldsymbol{b}^{\approx}$. Let C be the corresponding colorful choice for I that ray-embraces $\boldsymbol{b}$. Then, we set the predecessor of s to C. The properties of our perturbation ensure that we can compute C in polynomial time. Similarly, if s is a sink in G, we set its successor to the corresponding solution for the instance I. ☐

5 A Polynomial-Time Case

We show that for a special case of ColorfulCarathéodory, our formulation of Colorful-Carathéodory as a PPAD problem has algorithmic implications. Let $C_1, C_2 \in \mathbb{R}^d$ be two color classes and let $C \subseteq C_1 \cup C_2$ be a set. We call C an $(k, d-k)$-*colorful choice* for C_1 and C_2 if there are two subsets $C_1' \subseteq C_1$, $C_2' \subseteq C_2$ with $|C_1'| \leq k$ and $|C_2'| \leq d-k$. Now, given two color classes C_1, C_2 that each ray-embrace a point $\boldsymbol{b} \in \mathbb{R}^d$ and a number $k \in [d]_0$, we want to find an $(k, d-k)$-colorful choice that ray-embraces $\boldsymbol{b}$. It is a straightforward consequence of the colorful Carathéodory theorem that such a colorful choice always exists.

Using our techniques from Section 4, we present a weakly polynomial-time algorithm for this case. As described in Section 4.1, we construct implicitly a 1-dimensional polytopal complex, where at least one edge corresponds to a solution. Then, we apply binary search to find this edge. Since the length of the edges can be exponentially small in the length of the input, this results in a weakly polynomial-time algorithm.

Theorem 5.1. *Let $\boldsymbol{b} \in \mathbb{Q}^d$ be a point and let $C_1, C_2 \subset \mathbb{Q}^d$ be two sets of size d that ray-embrace $\boldsymbol{b}$. Furthermore, let $k \in [d-1]$ be a parameter. Then, there is an algorithm that computes a $(k, d-k)$-colorful choice C that ray-embraces $\mathbf{b}$ in weakly-polynomial time.*

For Sperner's lemma, it is well-known that a fully-labeled simplex can be found if there are only two labels by binary search. Essentially, this is also what the presented algorithm does: reducing the problem to Sperner's lemma and then applying binary search to find the right simplex. Since the computational problem Sperner is PPAD-complete even for $d = 2$, a polynomial-time generalization of this approach to three colors must use specific properties of the colorful Carathéodory instance under the assumption that no PPAD-complete problem can be solved in polynomial time.

6 Conclusion

We have shown that ColorfulCarathéodory lies in the intersection of PPAD and PLS. This also immediately implies that several illustrious problems associated with ColorfulCarathéodory, such as finding centerpoints or Tverberg partitions, belong to PPAD $\cap$ PLS.

Previously, the intersection PPAD $\cap$ PLS has been studied in the context of *continuous local search*: Daskalakis and Papadimitriou [10] define a subclass CLS $\subseteq$ PPAD $\cap$ PLS that "captures a particularly benign kind of local optimization". Daskalakis and Papadimitriou describe several interesting problems that lie in CLS but are not known to be solvable in polynomial time. Unfortunately, our results do not show that ColorfulCarathéodory lies in CLS, since we reduce ColorfulCarathéodory in d dimensions to Sperner in $d-1$ dimensions, and since Sperner is not known to be in CLS. Indeed, if Sperner's lemma could be shown to be in CLS, this would imply that PPAD $=$ CLS $\subseteq$ PLS, solving a major open problem. Thus, showing that Colorful-Carathéodory lies in CLS would require fundamentally new ideas, maybe exploiting the special structure of the resulting Sperner instance. On the other hand, it appears that Sperner is a more difficult problem than ColorfulCarathéodory, since Sperner is PPAD-complete for every fixed dimension larger than 1, whereas ColorfulCarathéodory becomes hard only in unbounded dimension. On the positive side, our perturbation results show that a polynomial-time algorithm for ColorfulCarathéodory, even under strong general position assumptions, would lead to polynomial-time algorithms for several well-studied problems in high-dimensional computational geometry.

Finally, it would also be interesting to find further special cases of ColorfulCarathéodory that are amenable to polynomial-time solutions. For example, can we extend our algorithm for two color classes to *three* color classes? We expect this to be difficult, due to an analogy between 1D-Sperner, which is in P, and 2D-Sperner, which is PPAD-complete. However, there seems to be no formal justification for this intuition.

References

[1] E. Aarts and J. K. Lenstra. *Local search in combinatorial optimization.* Princeton University Press, 2003.

[2] K. M. Anstreicher and R. A. Bosch. Long steps in an $O(n^3L)$ algorithm for linear programming. *Math. Program.*, 54(1-3):251–265, 1992.

[3] J. L. Arocha, I. Bárány, J. Bracho, R. Fabila, and L. Montejano. Very colorful theorems. *Discrete Comput. Geom.*, 42(2):142–154, 2009.

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

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

[6] S. Barman. Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Carathéodory's theorem. In *Proc. 47th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 361–369, 2015.

[7] V. Chvátal. *Linear programming.* W. H. Freeman and Company, 1983.

[8] D. I. A. Cohen. On the Sperner lemma. *J. Combinatorial Theory*, 2(4):585–587, 1967.

[9] G. B. Dantzig. *Linear Programming and Extensions.* Princeton University Press, 1963.

[10] C. Daskalakis and C. Papadimitriou. Continuous local search. In *Proc. 22nd Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 790–804, 2011.

[11] A. Holmsen. A combinatorial version of the colorful Carathéodory theorem. 2013.

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

[13] G. Kalai and R. Meshulam. A topological colorful Helly theorem. *Adv. Math.*, 191(2):305 – 311, 2005.

[14] S. Kapoor and P. M. Vaidya. Fast algorithms for convex quadratic programming and multicommodity flows. In *Proc. 18th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 147–159, 1986.

[15] M. K. Kozlov, S. P. Tarasov, and L. G. Khachiyan. The polynomial solvability of convex quadratic programming. *USSR Comput. Math. and Math. Phys.*, 20(5):223–228, 1980.

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

[17] J. Matoušek. *Using the Borsuk-Ulam theorem: lectures on topological methods in combinatorics and geometry*. Springer, 2008.

[18] J. Matousek and B. Gärtner. *Understanding and using linear programming*. Springer, 2007.

[19] N. Megiddo and R. Chandrasekaran. On the ϵ-perturbation method for avoiding degeneracy. *Oper. Res. Lett.*, 8(6):305–308, 1989.

[20] F. Meunier and P. Sarrabezolles. Colorful linear programming, Nash equilibrium, and pivots. arXiv:1409.3436, 2014.

[21] W. Michiels, E. Aarts, and J. Korst. *Theoretical aspects of local search*. Springer, 2007.

[22] W. Mulzer and Y. Stein. Computational aspects of the colorful Carathéodory theorem. In *Proc. 31st Int. Sympos. Comput. Geom. (SoCG)*, volume 34, pages 44–58, 2015.

[23] 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.

[24] R. Rado. A theorem on general measure. *J. London Math. Soc.*, 21:291–300 (1947), 1946.

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

[26] H. Tverberg. A generalization of Radon's theorem. *J. London Math. Soc*, 41(1):123–128, 1966.

[27] G. M. Ziegler. *Lectures on polytopes*. Springer, 1995.

Symbol	Definition
C_i	The ith color class. The d-set $C_i \subset \mathbb{R}^d$ ray-embraces $\boldsymbol{b}$.
A	The $(d \times d^2)$-matrix with C_1 as first d columns, C_2 as second d columns, and so on.
$c_{\boldsymbol{\mu}}$	The cost vector parameterized by a parameter vector $\boldsymbol{\mu} \in \mathbb{R}^d$. See (3).
L^{CC}; $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$	L^{CC} refers to the linear system $A\boldsymbol{x} = \boldsymbol{b}$, $\boldsymbol{x} \geq \boldsymbol{0}$ (see 2). $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$ denotes the linear program $\max c^T_{\boldsymbol{\mu}}\boldsymbol{x}$ s.t. L^{CC}.
$\mathcal{P}^{\mathrm{CC}}$	The polytope defined by L^{CC}.
f; $\operatorname{supp}(f)$; $\operatorname{ind}(B)$	For a face $f \subseteq \mathcal{P}^{\mathrm{CC}}$, we denote with $\operatorname{supp}(f)$ the indices of the columns in A that define it. For a set of columns B of A, we denote with $\operatorname{ind}(B)$ the indices of these columns.
$\Phi(f)$; $L^{\Phi}_{B,f}$	For a face f of $\mathcal{P}^{\mathrm{CC}}$, $\Phi(f)$ denotes the set of parameter vectors $\boldsymbol{\mu} \in \mathbb{R}^d$ such that f is optimal for $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$. The set $\Phi(f)$ can be described as the solution space to the linear system $L^{\Phi}_{B,f}$, where B is a feasible basis of a vertex of f.
$\mathcal{M}$	The set $\mathcal{M}$ contains all faces from the unit cube in $\mathbb{R}^d$ that set at least one coordinate to 1. Parameters from $\mathcal{M}$ control the colors of the defining columns of optimal faces (see Lemma 4.1).
$\mathcal{F}$	The set of faces f of $\mathcal{P}^{\mathrm{CC}}$ of that are optimal for some parameter vector in $\mathcal{M}$, i.e., the set of faces f with $\Phi(f) \cap \mathcal{M} \neq \emptyset$. $\mathcal{F}$ is a $(d-1)$-dimensional polyhedral complex.
$\mathcal{Q}$	The $(d-1)$-dimensional polytopal complex that consists of all elements $q = \Phi(f) \cap g$, where $f \in \mathcal{F}$ and g is a face of $\mathcal{M}$.
Δ; $\Delta_{[k]}$	Δ denotes the $(d-1)$-dimensional standard simplex and $\Delta_{[k]}$ denotes the face $\operatorname{conv}\{\boldsymbol{e}_i \mid i \in [k]\}$ of Δ.
$\mathcal{S}$	The set $\mathcal{S}$ contains the central projections of the faces of $\mathcal{M}$ onto Δ with the origin as center.
Φ_Δ; $\mathcal{Q}_\Delta$	$\Phi_\Delta(f)$ denotes the central projection of $\Phi(f) \cap \mathcal{M}$ onto Δ with center $\boldsymbol{0}$. The $(d-1)$-dimensional polytopal complex $\mathcal{Q}_\Delta$ consists of the projections of the elements in $\mathcal{Q}$ onto Δ. Each element q of $\mathcal{Q}_\Delta$ can be uniquely written as $q = \Phi_\Delta(f) \cap g$, where $f \in \mathcal{F}$ and $g \in \mathcal{S}$.
λ	The labeling function, see (5).
Σ; Σ_k; $\operatorname{enc}(\sigma)$	The set Σ_k, $k \in [d]$, consists of all $(k-1)$-simplices in $\operatorname{sd}\mathcal{Q}_\Delta$ that are contained in the face $\Delta_{[k]}$ of Δ. The set Σ is the union of all Σ_k. For a simplex $\sigma \in \Sigma$, we denote with $\operatorname{enc}(\sigma)$ its combinatorial encoding (see (6)).

Tab. 1: Notation reference.

A Polynomial-Time Reductions to the Colorful Carathéodory Problem

We begin by presenting the proofs of the centerpoint theorem, Tverberg's theorem, and the first selection lemma that use the colorful Carathéodory theorem. Afterwards, we show that these proofs can be interpreted as polynomial-time reductions to the corresponding computational problems.

Let $P \subset \mathbb{R}^d$ be a point set. We say a point $\boldsymbol{c} \in \mathbb{R}^d$ has *Tukey depth* τ with respect to P if and only if all closed halfspaces that contain $\boldsymbol{c}$ also contain at least τ points from P. The centerpoint theorem guarantees that there always exist points with large Tukey depth.

Theorem A.1 (Centerpoint theorem [24, Theorem 1])**.** *Let $P \subset \mathbb{R}^d$ be a point set. Then, there exists a point $\boldsymbol{q} \in \mathbb{R}^d$ with Tukey depth $\tau \geq \left\lceil \frac{|P|}{d+1} \right\rceil$.* □

We call a partition of P into m sets $T_1, \dots, T_m$ a *Tverberg m-partition* if and only if $\bigcap_{i=1}^m \operatorname{conv}(T_i) \neq \emptyset$. Tverberg's theorem guarantees that there are always large Tverberg partitions.

Theorem A.2 (Tverberg's theorem [26])**.** *Let $P \subset \mathbb{R}^d$ be a point set of size n. Then, there always exists a Tverberg $\left\lceil \frac{|P|}{d+1} \right\rceil$-partition for P. Equivalently, let P be of size $(m-1)(d+1)+1$ with $m \in \mathbb{N}$. Then, there exists a Tverberg m-partition for P.*

Note that Theorem A.2 directly implies Theorem A.1. A point $\boldsymbol{c}$ in the intersection of a Tverberg $\left\lceil \frac{|P|}{d+1} \right\rceil$-partition has Tukey depth at least $\left\lceil \frac{|P|}{d+1} \right\rceil$ since every halfspace that contains $\boldsymbol{c}$ must contain at least one point from each set in the Tverberg partition. We present Sarkaria's proof of Tverberg's theorem [25] with further simplifications by Bárány and Onn [5] and Arocha et al. [3]. The main tool is the following lemma that establishes a notion of duality between the intersection of convex hulls of low-dimensional point sets and the embrace of the origin of corresponding high-dimensional point sets. It was extracted from Sarkaria's proof by Arocha et al. [3]. In the following, we denote with $\otimes$ the tensor product.

In the following, we denote with $\otimes$ the binary function that maps two points $\boldsymbol{p} \in \mathbb{R}^d$, $\boldsymbol{q} \in \mathbb{R}^m$ to the point

$$\boldsymbol{p} \otimes \boldsymbol{q} = \begin{pmatrix} (\boldsymbol{q})_1 \boldsymbol{p} \\ (\boldsymbol{q})_2 \boldsymbol{p} \\ \vdots \\ (\boldsymbol{q})_m \boldsymbol{p} \end{pmatrix} \in \mathbb{R}^{dm}.$$

It is easy to verify that $\otimes$ is bilinear, i.e., for all $\boldsymbol{p}_1, \boldsymbol{p}_2 \in \mathbb{R}^d$, $\boldsymbol{q} \in \mathbb{R}^m$, and $\alpha_1, \alpha_2 \in \mathbb{R}$, we have

$$(\alpha_1 \boldsymbol{p}_1 + \alpha_2 \boldsymbol{p}_2) \otimes \boldsymbol{q} = \alpha_1 (\boldsymbol{p}_1 \otimes \boldsymbol{q}) + \alpha_2 (\boldsymbol{p}_2 \otimes \boldsymbol{q})$$

and similarly, for all $\boldsymbol{p} \in \mathbb{R}^d$, $\boldsymbol{q}_1, \boldsymbol{q}_2 \in \mathbb{R}^m$, and $\alpha_1, \alpha_2 \in \mathbb{R}$, we have

$$\boldsymbol{p} \otimes (\alpha_1 \boldsymbol{q}_1 + \alpha_2 \boldsymbol{q}_2) = \alpha_1 (\boldsymbol{p} \otimes \boldsymbol{q}_1) + \alpha_2 (\boldsymbol{p} \otimes \boldsymbol{q}_2).$$

Lemma A.3 (Sarkaria's lemma [25], [3, Lemma 2])**.** *Let $P_1, \dots, P_m \subset \mathbb{R}^d$ be m point sets and let $\boldsymbol{q}_1, \dots, \boldsymbol{q}_m \subset \mathbb{R}^{m-1}$ be m vectors with $\boldsymbol{q}_i = \boldsymbol{e}_i$ for $i \in [m-1]$ and $\boldsymbol{q}_m = -\boldsymbol{1}$. For $i \in [m]$, we define*

$$\widehat{P}_i = \left\{ \begin{pmatrix} \boldsymbol{p} \\ 1 \end{pmatrix} \otimes \boldsymbol{q}_i \,\middle|\, \boldsymbol{p} \in P_i \right\} \subset \mathbb{R}^{(d+1)(m-1)}.$$

Then, the intersection of convex hulls $\bigcap_{i=1}^m \operatorname{conv}(P_i)$ is nonempty if and only if $\bigcup_{i=1}^m \widehat{P}_i$ embraces the origin.

Proof. Assume there is a point $\boldsymbol{p}^\star \in \bigcap_{i=1}^m \operatorname{conv}(P_i)$. For $i \in [m]$ and $\boldsymbol{p} \in P_i$, there then exist coefficients $\lambda_{i,\boldsymbol{p}} \in \mathbb{R}_+$ that sum to 1 such that $\boldsymbol{p}^\star = \sum_{\boldsymbol{p}\in P_i} \lambda_{i,\boldsymbol{p}}$. Consider the points $\hat{\boldsymbol{p}}_i \in \operatorname{conv}\left(\widehat{P}_i\right)$, $i \in [m]$, that we obtain by using the same convex coefficients for the points in $\widehat{P}_i$, i.e., set

$$\hat{\boldsymbol{p}}_i = \sum_{\boldsymbol{p}\in P_i} \lambda_{i,\boldsymbol{p}} \left(\begin{pmatrix}\boldsymbol{p}\\1\end{pmatrix} \otimes \boldsymbol{q}_i\right) \in \operatorname{conv}\left(\widehat{P}_i\right).$$

We claim that $\sum_{i=1}^m \hat{\boldsymbol{p}}_i = \mathbf{0}$ and thus $\mathbf{0} \in \operatorname{conv}\left(\bigcup_{i=1}^m \widehat{P}_i\right)$. Indeed, we have

$$\begin{aligned}\sum_{i=1}^m \hat{\boldsymbol{p}}_i = \sum_{i=1}^m \sum_{\boldsymbol{p}\in P_i} \lambda_{i,\boldsymbol{p}} \left(\begin{pmatrix}\boldsymbol{p}\\1\end{pmatrix} \otimes \boldsymbol{q}_i\right) = \sum_{i=1}^m \left(\sum_{\boldsymbol{p}\in P_i} \lambda_{i,\boldsymbol{p}} \begin{pmatrix}\boldsymbol{p}\\1\end{pmatrix}\right) \otimes \boldsymbol{q}_i = \sum_{i=1}^m \begin{pmatrix}\boldsymbol{p}^\star\\1\end{pmatrix} \otimes \boldsymbol{q}_i \\ = \begin{pmatrix}\boldsymbol{p}^\star\\1\end{pmatrix} \otimes \left(\sum_{i=1}^m \boldsymbol{q}_i\right) = \begin{pmatrix}\boldsymbol{p}^\star\\1\end{pmatrix} \otimes \mathbf{0} = \mathbf{0},\end{aligned}$$

where we use the fact that $\otimes$ is bilinear.

Assume now that $\bigcup_{i=1}^m \widehat{P}_i$ embraces the origin and we want to show that $\bigcap_{i=1}^m \operatorname{conv}(P_i)$ is nonempty. Then, we can express the origin as a convex combination $\sum_{i=1}^m \sum_{\hat{\boldsymbol{p}}\in\widehat{P}_i} \lambda_{i,\hat{\boldsymbol{p}}}\hat{\boldsymbol{p}}$ with $\lambda_{i,\hat{\boldsymbol{p}}} \in \mathbb{R}_+$ for $i \in [m]$ and $\hat{\boldsymbol{p}} \in \widehat{P}_i$, and $\sum_{i=1}^m \sum_{\hat{\boldsymbol{p}}\in\widehat{P}_i} \lambda_{i,\hat{\boldsymbol{p}}} = 1$. Hence, we have

$$\mathbf{0} = \sum_{i=1}^m \sum_{\hat{\boldsymbol{p}}\in\widehat{P}_i} \lambda_{i,\hat{\boldsymbol{p}}} \left(\begin{pmatrix}\boldsymbol{p}\\1\end{pmatrix} \otimes \boldsymbol{q}_i\right) = \sum_{i=1}^m \left(\sum_{\hat{\boldsymbol{p}}\in\widehat{P}_i} \lambda_{i,\hat{\boldsymbol{p}}} \begin{pmatrix}\boldsymbol{p}\\1\end{pmatrix}\right) \otimes \boldsymbol{q}_i,$$

where we use again the fact that $\otimes$ is bilinear. By the choice of $\boldsymbol{q}_1, \dots, \boldsymbol{q}_m$, there is (up to multiplication with a scalar) exactly one linear dependency: $\mathbf{0} = \sum_{i=1}^m \boldsymbol{q}_i$. Thus,

$$\sum_{\hat{\boldsymbol{p}}\in\widehat{P}_1} \lambda_{1,\hat{\boldsymbol{p}}} \begin{pmatrix}\boldsymbol{p}\\1\end{pmatrix} = \cdots = \sum_{\hat{\boldsymbol{p}}\in\widehat{P}_m} \lambda_{m,\hat{\boldsymbol{p}}} \begin{pmatrix}\boldsymbol{p}\\1\end{pmatrix} = \begin{pmatrix}\boldsymbol{p}^\star\\c\end{pmatrix},$$

where $\boldsymbol{p}^\star \in \mathbb{R}^d$ and $c \in \mathbb{R}$. In particular, the last equality implies that

$$\sum_{\hat{\boldsymbol{p}}\in\widehat{P}_1} \lambda_{1,\hat{\boldsymbol{p}}} = \cdots = \sum_{\hat{\boldsymbol{p}}\in\widehat{P}_m} \lambda_{m,\hat{\boldsymbol{p}}} = c.$$

Now, since for all $i \in [m]$ and $\hat{\boldsymbol{p}} \in \widehat{P}_i$, the coefficient $\lambda_{i,\hat{\boldsymbol{p}}}$ is nonnegative and since the sum $\sum_{i\in[m]} \sum_{\hat{\boldsymbol{p}}\in\widehat{P}_i} \lambda_{i,\hat{\boldsymbol{p}}}$ is 1, we must have $c = 1/m \in (0, 1]$. Hence, the point $m\boldsymbol{p}^\star$ is common to all convex hulls $\operatorname{conv}(P_1), \dots, \operatorname{conv}(P_m)$. □

Little work is now left to obtain Tverberg's theorem from Lemma A.3 and the colorful Carathéodory theorem.

Proof of Theorem A.2. Let $P = \{\boldsymbol{p}_1, \dots, \boldsymbol{p}_n\} \subset \mathbb{R}^d$ be a point set of size $n = (d+1)(m-1)+1$ and let $P_1, \dots, P_m$ denote m copies of P. For each set $P_j \subset \mathbb{R}^d$, $j \in [m]$, we construct a $((d+1)(m-1))$-dimensional set $\widehat{P}_j$ as in Lemma A.3, i.e.,

$$\widehat{P}_j = \left\{\hat{\boldsymbol{p}}_{i,j} = \boldsymbol{p}_i \otimes \boldsymbol{q}_j \,\middle|\, \boldsymbol{p}_i \in P\right\} \subset \mathbb{R}^{(d+1)(m-1)} = \mathbb{R}^{n-1}.$$

For $i \in [n]$, we denote with $\widehat{C}_i \subseteq \bigcup_{j=1}^m \widehat{P}_j$ the set of points $\left\{\hat{\boldsymbol{p}}_{i,j} \,\middle|\, j \in [m]\right\}$ that correspond to $\boldsymbol{p}_i \in P$ and we color these points with color i. For $i \in [n]$, note that Lemma A.3 applied to m copies of the singleton set $\{\boldsymbol{p}_i\} \subseteq P$ guarantees that the color class $\widehat{C}_i \in \mathbb{R}^{n-1}$ embraces the origin. Hence, we have n color classes $\widehat{C}_1, \ldots, \widehat{C}_n$ that embrace the origin in $\mathbb{R}^{n-1}$. Now, by Theorem 1.1, there is a colorful choice $\widehat{C} = \{\hat{\boldsymbol{c}}_1, \ldots, \hat{\boldsymbol{c}}_n\} \subseteq \bigcup_{i=1}^n \widehat{C}_i$ with $\hat{\boldsymbol{c}}_i \in \widehat{C}_i$ that embraces the origin, too. Because $\widehat{C}$ embraces the origin, Lemma A.3 guarantees that the convex hulls of the sets $T_j = \left\{\boldsymbol{p}_i \in P \,\middle|\, \hat{\boldsymbol{p}}_{i,j} \in \widehat{C}\right\}$, $j \in [m]$, have a point in common. Moreover, since all points in $\bigcup_{j=1}^m \widehat{P}_j$ that correspond to the same point in P have the same color, each point $\boldsymbol{p}_i \in P$ appears in exactly one set T_j, $j \in [m]$. Thus, $\mathcal{T} = \{T_1, \ldots, T_j\}$ is a Tverberg m-partition of P. ☐

Similar to the Tukey depth, the simplicial depth is a further notion of data depth. Let again be $P \subset \mathbb{R}^d$ be a point set and $\boldsymbol{q} \in \mathbb{R}^d$ a point. Then, the *simplicial depth* $\delta_P(\boldsymbol{q})$ of $\boldsymbol{q}$ with respect to P is the number of distinct d-simplices that contain $\boldsymbol{q}$ with vertices in P. The first selection lemma states that for fixed d, there is always a point with asymptotic optimal simplicial depth.

Theorem A.4 (First selection lemma [4, Theorem 5.1])**.** *Let $P \subset \mathbb{R}^d$ be a set of points and consider d constant. Then, there exists a point $\boldsymbol{q} \in \mathbb{R}^d$ with $\delta_P(\boldsymbol{q}) = \Omega\left(|P|^{d+1}\right)$.*

The main argument of Bárány's proof of the first selection lemma is the following lemma.

Lemma A.5. *Let $P \subset \mathbb{R}^d$ be a point set and let $\mathcal{T}$ be a Tverberg m-partition of P, where $m \in \mathbb{N}$. Then any point $\boldsymbol{c} \in \bigcap_{T \in \mathcal{T}} \operatorname{conv}(T)$ has simplicial depth $\sigma_P(\boldsymbol{c})$ at least $\left\lceil \frac{m^{d+1}}{(d+1)^{d+1}} \right\rceil$.*

Proof. Let T_i denote the ith element of $\mathcal{T}$ and color it with color i. Now by Theorem 1.1, there exists for every $(d+1)$-subset $I \subseteq [m]$ a colorful choice C_I with respect to the color classes T_i, $i \in I$, that embraces $\boldsymbol{c}$. Furthermore, each index set I induces a unique colorful choice C_I. Thus, there are at least $\binom{m}{d+1} \geq \frac{m^{d+1}}{(d+1)^{d+1}}$ distinct $\boldsymbol{c}$-embracing d-simplices with vertices in P. ☐

The first selection lemma is now an immediate consequence of Lemma A.5 and Theorem A.2.

We define the computational problems that correspond to the centerpoint theorem, Tverberg's theorem, and the first selection lemma as follows.

Definition A.6. We define the following search problems:

- CENTERPOINT

  Given a set $P \subset \mathbb{Q}^d$ of size n,

  Find a centerpoint.

- TVERBERG

  Given a set $P \subset \mathbb{Q}^d$ of size n,

  Find a Tverberg $\lceil \frac{n}{d+1} \rceil$-partition.

- SIMPLICIALCENTER

  Given a set $P \subset \mathbb{Q}^d$ of size n,

  Find a point $\boldsymbol{q} \in \mathbb{Q}^d$ with $\sigma_P(\boldsymbol{q}) \geq f(d)n^{d+1}$, where $f : \mathbb{N} \mapsto \mathbb{R}_+$ is an arbitrary but fixed function.

Finally, interpreting the presented proofs as algorithms, we obtain the following result.

Lemma A.7. *Given access to an oracle for* ColorfulCarathéodory*,* Tverberg *can be solved in* $O(n^3)$ *time. Furthermore,* Centerpoint *and* SimplicialCenter *can be solved in* $O(Ln^3)$ *time, where* L *is the length of the input.*

Proof. As show in the proof of Theorem A.2, to compute a Tverberg partition, it suffices to lift $m = \left\lceil \frac{n}{d+1} \right\rceil$ copies of the input point set $P \subset \mathbb{Q}^d$ with Lemma A.3 and then query the oracle for ColorfulCarathéodory. Lifting one point needs $O(dm) = O(n)$ time and hence we need $O(n^3)$ time in total. Then, any point in the intersection of the computed Tverberg m-partition $\mathcal{T} = \{T_1, \dots, T_m\}$ is a solution to Centerpoint and SimplicialCenter. Using the algorithm from [2], we can compute a Tverberg point in time $O(Ln^3)$ by solving the linear program

$$\left(\begin{array}{ccc|c} \boxed{\begin{matrix} & T_1 & \\ 1 & \dots & 1\end{matrix}} & 0 & & \begin{matrix}-1 & & \\ & \ddots & \\ & & -1 \\ 0 & \dots & 0\end{matrix} \\ \cline{4-4} & \ddots & & \vdots \\ \cline{4-4} 0 & & \boxed{\begin{matrix} & T_m & \\ 1 & \dots & 1\end{matrix}} & \begin{matrix}-1&&\\&\ddots&\\&&-1\\0&\dots&0\end{matrix} \end{array}\right) \boldsymbol{x} = \left(\begin{array}{c} 0\\ \vdots \\ 0 \\ 1 \\ \hline \vdots \\ \hline 0 \\ \vdots \\ 0 \\ 1 \end{array}\right) \quad \text{s.t. } \boldsymbol{x} \geq \boldsymbol{0},$$

where L is the length of the input. ☐

B Equivalent Instances of the Colorful Carathéodory Problem in General Position

The application of Sarkaria's lemma in the reductions to ColorfulCarathéodory creates color classes whose positive span does not have full dimension. To be able to transfer upper bounds on the complexity of ColorfulCarathéodory to its descendants, we need to be able to deal with degenerate position. In this chapter, we show how to ensure general position of Colorful-Carathéodory instances by extending known perturbation techniques for linear programming to our setting. More formally, let $I = (C_1, \dots, C_d, \boldsymbol{b})$ be a ColorfulCarathéodory instance, where $\boldsymbol{b} \in \mathbb{Q}^d \setminus \{\boldsymbol{0}\}$ and each color class $C_i \subset \mathbb{Q}^d$, $i \in [d]$, ray-embraces $\boldsymbol{b}$. Then, we want to construct in polynomial time d sets $C_1^{\approx}, \dots, C_d^{\approx} \subset \mathbb{Z}^d$ and a point $\boldsymbol{b}^{\approx} \in \mathbb{Z}^d$ that have the following properties:

(P1) **Valid instance with integer coordinates:** The points $\{\boldsymbol{b}^{\approx}\} \cup \left(\bigcup_{i=1}^d C_i^{\approx}\right) \subset \mathbb{Z}^d$ have integer coordinates. Furthermore, the point $\boldsymbol{b}^{\approx}$ is not the origin and each color class $C_i^{\approx}$, $i \in [d]$, ray-embraces $\boldsymbol{b}^{\approx}$ and has size d.

(P2) **$\boldsymbol{b}$ avoids linear subspaces:** The point $\boldsymbol{b}^{\approx}$ is not contained in the linear span of any $(d-1)$-subset of $\bigcup_{i=1}^d C_i^{\approx}$.

(P3) **Polynomial-time equivalent solutions:** Given a colorful choice $C^{\approx} \subseteq \bigcup_{i=1}^d C_i^{\approx}$ that ray-embraces $\boldsymbol{b}^{\approx}$, we can compute in polynomial time a colorful choice $C \subseteq \bigcup_{i=1}^d C_i$ that ray-embraces $\boldsymbol{b}$.

Note that by (P2), if $P \subset \bigcup_{i=1}^d C_i^{\approx}$ ray-embraces $\boldsymbol{b}^{\approx}$, then $|P| \geq d$ and thus $\boldsymbol{b}^{\approx} \in \operatorname{int} \operatorname{pos}(P)$. In particular by (P1), $\boldsymbol{b}^{\approx}$ is contained in the interior of $\operatorname{pos}(C_i^{\approx})$ for $i \in [d]$.

In the next section, we develop tools to ensure non-degeneracy of linear systems by a small deterministic perturbation of polynomial bit-complexity. The approach is similar to already existing perturbation techniques for linear programming as in [9, Section 10-2] and [19] but extends to a more general setting in which the matrix is also perturbed. Based on these results, we then show in Section B.2 how to construct ColorfulCarathéodory instances with properties (P1)–(P3).

B.1 Polynomials with Bounded Integer Coefficients

In the following, we consider equation systems

$$L_\varepsilon : A\boldsymbol{x} = \boldsymbol{b}, \tag{7}$$

where A is a $(d \times n)$-matrix with $n \geq d$ and $\boldsymbol{b}$ is a d-dimensional vector. Furthermore, the entries of both A and $\boldsymbol{b}$ are polynomials in ε with integer coefficients. For a fixed $\tau \in \mathbb{R}$, we denote with $A(\tau)$ and $\boldsymbol{b}(\tau)$ the matrix A and the vector $\boldsymbol{b}$ that we obtain by setting ε to τ in A and $\boldsymbol{b}$, respectively. Similarly, we denote with L_τ the linear system $L_\tau : A(\tau)\boldsymbol{x} = \boldsymbol{b}(\tau)$. We show that for any fixed $\tau > 0$ that is sufficiently small in the size of the coefficients in the polynomials, the linear system L_τ is non-degenerate.

For $m \in \mathbb{N}$, we denote with

$$\mathbb{P}[m] = \left\{ p(\varepsilon) = \sum_{i=0}^{k} \alpha_i \varepsilon^i \,\middle|\, k \in \mathbb{N}_0, \text{ and } |\alpha_i| \in [m]_0 \text{ for } i \in [k]_0 \right\}$$

the set of polynomials with integer coefficients that have absolute value at most m. The following lemma guarantees that no polynomial in $\mathbb{P}[m]$ has a root in a specific interval whose length is inverse proportional to m.

Lemma B.1. *Let $p \in \mathbb{P}[m]$ be a nontrivial polynomial with $m \in \mathbb{N}$. Then, for all $\varepsilon \in \left(0, \frac{1}{2m}\right)$, we have $p(\varepsilon) \neq 0$.*

Proof. We write $p(\varepsilon) = \sum_{i=0}^{k} \alpha_i \varepsilon^i$. Let $j = \min\{i \in [k]_0 \mid \alpha_i \neq 0\}$. Since p is nontrivial, j exists. Without loss of generality, we assume $\alpha_j > 0$ (otherwise, we multiply $p(\varepsilon)$ by -1). For all $\varepsilon \in \left(0, \frac{1}{2m}\right)$, we have

$$p(\varepsilon) = \sum_{i=0}^{k} \alpha_i \varepsilon^i \geq \varepsilon^j - 2m\varepsilon^{j+1} = \varepsilon^j \left(1 - 2m\varepsilon\right) > 0$$

since $\varepsilon < \frac{1}{2m}$ and hence $p(\varepsilon) \neq 0$ for all $\varepsilon \in \left(0, \frac{1}{2m}\right)$. ☐

We now use Lemma B.1 to prove non-degeneracy of the linear system L_ε if ε is fixed but small enough and the degrees of the monomials in L_ε are sufficiently separated. We say d polynomials $p_1, \ldots, p_d \in \mathbb{P}[m]$ are $(k_1, \ldots, k_d)$*-separated with gap g* if p_i has a nontrivial monomial of degree k_i and p_i has no nontrivial monomial of a degree in $\{k_j - g, \ldots, k_j + g \mid j \in [d] \setminus \{i\}\} \cup \{k_i - g, \ldots, k_i - 1\}$.

Lemma B.2. *Let $L_\varepsilon : A\boldsymbol{x} = \boldsymbol{b}$ be a system of equations as defined in (7) such that the entries of A and $\boldsymbol{b}$ are polynomials in $\mathbb{P}[m]$, where $m \in \mathbb{N}$. Furthermore, suppose that the polynomials in A have degree at most k_0 and $(\boldsymbol{b})_1, \ldots, (\boldsymbol{b})_d$ are $(k_1, \ldots, k_d)$-separated with gap $(d-1)k_0$. Set*

$$M = d!(k_0 + 1)^{d-1}(k + 1)m^d,$$

where k is the maximum degree of $(\boldsymbol{b})_1, \ldots, (\boldsymbol{b})_d$. Then, for all $\varepsilon \in \left(0, \frac{1}{2M}\right)$, the linear system L_ε is non-degenerate.

Proof. We show that for all fixed $\tau \in \left(0, \frac{1}{2M}\right)$, the vector $\boldsymbol{b}(\tau)$ is not contained in the linear span of any $d-1$ columns from $A(\tau)$. We can ensure that $A(\tau)$ has rank d for all fixed $\tau \geq 0$ by extending A with the canonical basis of $\mathbb{R}^d$. Then, the entries of the extended matrix are still polynomials from $\mathbb{P}[m]$ and their degrees are at most k_0. Moreover, if for some fixed $\tau \in \left(0, \frac{1}{2M}\right)$, there are $d-1$ columns from the original matrix whose linear span contains $\boldsymbol{b}(\tau)$, then the same holds for the extended matrix.

Let now $\tau \in \left(0, \frac{1}{2M}\right)$ be fixed and let A' be a submatrix of A such that $A'(\tau)$ is a basis of $A(\tau)$. Then, the linear system

$$L' : A'(\tau)\boldsymbol{x} = \boldsymbol{b}(\tau)$$

has a unique solution $\boldsymbol{x}^\star$. By Cramer's rule, we have

$$(\boldsymbol{x}^\star)_j = \frac{\det A'_j(\tau)}{\det A'(\tau)},$$

where $j \in [d]$ and A'_j is obtained from the matrix A' by replacing the jth column with $\boldsymbol{b}$. Using Laplace expansion, we can express $\det A'_j$ as

$$\det A'_j = \sum_{i=1}^{d} (-1)^{i+j} b_i \det C_{i,j},$$

where $b_i = (\boldsymbol{b})_i$ and $C_{i,j}$ is the matrix that we obtain by omitting the ith row and the jth column from A'_j. Next, we apply the Leibniz formula and write $\det C_{i,j}$ as

$$\det C_{i,j} = \sum_{\sigma \in S_{d-1}} \operatorname{sgn}(\sigma) \prod_{i=1}^{d-1} (C_{i,j})_{i,\sigma(i)} = c_{i,j}(\varepsilon),$$

where $c_{i,j}(\varepsilon)$ is a polynomial in ε. Since the polynomials in A' have degree at most k_0, the degree of $c_{i,j}$ is at most $(d-1)k_0$. Because the polynomials in A' have integer coefficients with absolute value at most m, the coefficients of $c_{i,j}$ are integers, and the sum of their absolute values can be bounded by $M' = (d-1)!\big((k_0+1)m\big)^{d-1}$. Hence, $c_{i,j} \in \mathbb{P}[M']$. Now, since $\det A'(\tau) \neq 0$, at least one of the polynomials $c_{1,j}, \ldots, c_{d,j}$, say $c_{i^\star,j}$, is nontrivial. Let $k'_{i^\star} \leq (d-1)k_0$ be the minimum degree of a nontrivial monomial in $c_{i^\star,j}$. First, we observe that since $b_{i^\star}$ has a nontrivial monomial of degree $k_{i^\star}$ and no nontrivial monomial of degree $k_{i^\star} - (d-1)k_0, \ldots, k_{i^\star} - 1$, the polynomial $(-1)^{i^\star+j} b_{i^\star} c_{i^\star,j}$ has a nontrivial monomial of degree $k' = k_{i^\star} + k'_{i^\star}$. Second, for $i \in [d]$, $i \neq i^\star$, the polynomial $(-1)^{i+j} b_i c_{i,j}$ has no monomial of degree k' since $c_{i,j}$ has degree at most $(d-1)k_0$ and the polynomials $b_1, \ldots, b_d$ are $(k_1, \ldots, k_d)$-separated with gap $(d-1)k_0$. Thus, $\det A'_j$ is a nontrivial polynomial. Moreover, since the polynomials b_i and $c_{i,j}$ have integer coefficients for $i \in [d]$, so does $\det A'_j$. Using that the sum of absolute values of the coefficients of $c_{i,j}$ is bounded by M', we can bound the sum of absolute values of coefficients in $\det A'_j$ by $M = d(k+1)mM'$ and hence $\det A'_j \in \mathbb{P}[M]$, where $k = \max\{\deg b_i \mid i \in [d]\}$. Then, Lemma B.1 guarantees that $\det A'_j$ has no root in the interval $\left(0, \frac{1}{2M}\right)$. In particular, $\det A'_j(\tau) \neq 0$ and hence $(\boldsymbol{x}^\star)_j \neq 0$ for all $j \in [d]$. This means that $\boldsymbol{b}(\tau)$ is not contained in the linear span of any $d-1$ columns from $A(\tau)$. Since $\tau \in \left(0, \frac{1}{2M}\right)$ was arbitrary, the claim follows. ☐

B.2 Construction

Let $C'_1, \ldots, C'_d \subset \mathbb{Q}^d$ be d sets that ray-embrace $\boldsymbol{b}' \in \mathbb{Q}^d$. By applying Carathéodory's theorem, we can ensure that $|C'_i| \leq d$ for $i \in [d]$. First, we rescale the points to the integer grid. For a point

$\boldsymbol{p}' \in \mathbb{Q}_d$, we set $z(\boldsymbol{p}') = |\psi|\boldsymbol{p}'$, where $\psi \in \mathbb{Z}$ is the absolute value of the least common multiple of the denominators of $(\boldsymbol{p}')_1, \dots, (\boldsymbol{p}')_d$. Clearly, $z(\boldsymbol{p}')$ has integer coordinates and can be represented with a number of bits polynomial in the number of bits needed for $\boldsymbol{p}'$. For $i \in [d]$, let $C_i = \{z(\boldsymbol{p}') \mid \boldsymbol{p}' \in C_i'\}$ be the rescaling of C_i', and set $\boldsymbol{b} = z(\boldsymbol{b}')$. Then, the bit complexity of the ColorfulCarathéodory instance $C_1, \dots, C_d, \boldsymbol{b}$ is polynomial in the bit-complexity of the original instance. Moreover, since $\operatorname{pos}(\boldsymbol{p}') = \operatorname{pos}(z(\boldsymbol{p}'))$ for all $\boldsymbol{p}' \in \mathbb{Q}^d$, the rescaled color classes C_i, $i \in [d]$, ray-embrace $\boldsymbol{b}$ and if a colorful choice $C \subseteq \bigcup_{i=1}^d C_i$ ray-embraces $\boldsymbol{b}$, then the original points $C' \subset \bigcup_{i=1}^d C_i'$ ray-embrace $\boldsymbol{b}'$. By a similar rescaling, we can further assume that $\|\boldsymbol{b}\|_1 \geq \|\boldsymbol{p}\|_1$ for all $\boldsymbol{p} \in \bigcup_{i=1}^d C_i$.

We now sketch how the remaining construction of the equivalent instance $C_1^{\approx}, \dots, C_d^{\approx}, \boldsymbol{b}^{\approx}$ in general position proceeds. First, we ensure for $i \in [d]$ that $\boldsymbol{b}$ lies in the interior of $\operatorname{pos}(C_i)$ by replacing each point $\boldsymbol{p}$ in C_i by a set $P_\varepsilon(\boldsymbol{p})$ of slightly perturbed points that contain $\boldsymbol{p}$ in the interior of their convex hull. Second, we perturb $\boldsymbol{b}$. Lemma B.2 then shows that in both steps a perturbation of polynomial bit-complexity suffices to ensure properties (P2) and (P3).

For a point $\boldsymbol{p} \in \mathbb{R}^d$, we denote with

$$P_\varepsilon(\boldsymbol{p}) = \{\boldsymbol{p} + \varepsilon \boldsymbol{e}_i, \boldsymbol{p} - \varepsilon \boldsymbol{e}_i \mid i \in [d]\}$$

the vertices of the ℓ_1-sphere around $\boldsymbol{p}$ with radius ε. Let $C_i(\varepsilon) = \bigcup_{\boldsymbol{p} \in C_i} P_\varepsilon(\boldsymbol{p})$, $i \in [d]$, denote the ith color class in which all points $\boldsymbol{p}$ have been replaced by the corresponding set $P_\varepsilon(\boldsymbol{p})$. Since for $i \in [d]$, we have $\boldsymbol{b} \in \operatorname{pos}(C_i)$ and since each point $\boldsymbol{p} \in C_i$ is contained in the interior of $\operatorname{pos}(P_\varepsilon(\boldsymbol{p}))$, it follows that $\boldsymbol{b} \in \operatorname{int}\operatorname{pos}(C_i(\varepsilon))$ for $\varepsilon > 0$. Next, we denote with

$$\boldsymbol{b}(\varepsilon) = \boldsymbol{b} + \begin{pmatrix} \varepsilon^d \\ \varepsilon^{2d} \\ \vdots \\ \varepsilon^{d^2} \end{pmatrix} \in \mathbb{R}^d$$

the vector $\boldsymbol{b}$ that is perturbed by a vector from the moment curve. The following lemma shows that for ε small enough, Property (P2) holds for $C_1(\varepsilon), \dots, C_d(\varepsilon)$ and $\boldsymbol{b}(\varepsilon)$. Let m be the largest absolute value of a coordinate in $C_1, \dots, C_d, \boldsymbol{b}$ and set $N = d! m^d$.

Lemma B.3. *For all $\varepsilon \in (0, N^{-2}]$, there is no $(d-1)$-subset $P \subset \bigcup_{i=1}^d C_i(\varepsilon)$ with $\boldsymbol{b}(\varepsilon) \in \operatorname{span} P$.*

Proof. Let A denote the matrix $\big(C_1(\varepsilon) \dots C_d(\varepsilon)\big)$. Then, there exists a subset $P \subset \bigcup_{i=1}^d C_i(\varepsilon)$ with $|P| < d$ that contains $\boldsymbol{b}(\varepsilon)$ in its linear span if and only if the linear system $L_\varepsilon : A\boldsymbol{x} = \boldsymbol{b}(\varepsilon)$ is degenerate. The polynomials in A all have degree at most 1 and the polynomials $(\boldsymbol{b}(\varepsilon))_i$, $i \in [d]$, are $(d, 2d, \dots, d^2)$-separated with gap $d-1$. Setting $k_0 = 1$ and $k = d^2$ in Lemma B.2 implies that L_ε is non-degenerate for all $\varepsilon \in \left(0, \frac{1}{2M}\right)$, where $M = d! 2^{d-1}(d^2+1)m^d$. Assuming that $m \geq 2$ and that $d \geq 4$, we can upper bound 2^d by m^d and (d^2+1) by $d!$. Hence, we have

$$2M = d! 2^d (d^2+1) m^d < \left(d! m^d\right)^2 = N^2,$$

and thus the claim follows. ☐

In the following, we set ε_0 to N^{-2}. Note that Lemma B.3 holds in particular for $\varepsilon = \varepsilon_0$, and thus a deterministic perturbation of polynomial bit-complexity suffices. In the next lemma, we show that the perturbed color classes still ray-embrace the perturbed $\boldsymbol{b}$.

Lemma B.4. *For $i \in [d]$, the set $C_i(\varepsilon_0)$ ray-embraces $\boldsymbol{b}(\varepsilon_0)$.*

Proof. Fix some color class C_i and let $\boldsymbol{m}_{\varepsilon_0} = \boldsymbol{b}(\varepsilon) - \boldsymbol{b}$ be the perturbation vector for $\boldsymbol{b}$. Since C_i ray-embraces $\boldsymbol{b}$, we can express $\boldsymbol{b}$ as a positive combination $\sum_{\boldsymbol{p}\in C_i} \psi_{\boldsymbol{p}}\boldsymbol{p}$, where $\psi_{\boldsymbol{p}} \geq 0$ for all $\boldsymbol{p} \in C_i$. Then,

$$\boldsymbol{b}(\varepsilon_0) = \boldsymbol{b} + \boldsymbol{m}_{\varepsilon_0} = \left(\sum_{\boldsymbol{p}\in C_i} \psi_{\boldsymbol{p}}\boldsymbol{p}\right) + \boldsymbol{m}_{\varepsilon_0} = \sum_{\boldsymbol{p}\in C_i} \psi_{\boldsymbol{p}} \left(\boldsymbol{p} + \frac{1}{s}\boldsymbol{m}_{\varepsilon_0}\right),$$

where $s = \sum_{\boldsymbol{p}\in C_i} \psi_{\boldsymbol{p}}$. We show that $\boldsymbol{p} + \frac{1}{s}\boldsymbol{m}_{\varepsilon_0} \in \operatorname{pos}(P_{\varepsilon_0}(\boldsymbol{p}))$ for all $\boldsymbol{p} \in C_i$. Since $P_{\varepsilon_0}(\boldsymbol{p}) \subseteq C_i(\varepsilon_0)$ for all $\boldsymbol{p} \in C_i$, this then implies $\boldsymbol{b}(\varepsilon_0) \in \operatorname{pos}(C_i(\varepsilon_0))$. First, we claim that $s \geq 1$. Indeed, we have

$$\|\boldsymbol{b}\|_1 = \left\|\sum_{\boldsymbol{p}\in C_i} \psi_{\boldsymbol{p}}\boldsymbol{p}\right\|_1 \leq \sum_{\boldsymbol{p}\in C_i} \psi_{\boldsymbol{p}} \|\boldsymbol{p}\|_1 \leq s\|\boldsymbol{b}\|_1,$$

where the last inequality is due to our assumption $\|\boldsymbol{b}\|_1 \geq \|\boldsymbol{p}\|_1$, for $\boldsymbol{p} \in C_i$. Now,

$$\left\|\frac{1}{s}\boldsymbol{m}_{\varepsilon_0}\right\|_1 < d\varepsilon_0^d \leq \varepsilon_0,$$

for $\varepsilon_0 \leq 1/2$, and thus $\boldsymbol{p} + \frac{1}{s}\boldsymbol{m}_{\varepsilon_0}$ lies in the ℓ_1-sphere around $\boldsymbol{p}$ with radius ε_0 for all $\boldsymbol{p} \in C_i$. By construction of $P_{\varepsilon_0}(\boldsymbol{p})$, we then have $\boldsymbol{p} + \frac{1}{s}\boldsymbol{m}_{\varepsilon_0} \in \operatorname{conv}(P_{\varepsilon_0}(\boldsymbol{p})) \subset \operatorname{pos}(P_{\varepsilon_0}(\boldsymbol{p}))$, as claimed. □

As a consequence of Lemma B.3, we can show that colorful choices for the perturbed instance that ray-embrace $\boldsymbol{b}(\varepsilon_0)$, ray-embrace $\boldsymbol{b}$ if the perturbation is removed.

Lemma B.5. *Let $C = \{\boldsymbol{c}_1, \ldots, \boldsymbol{c}_d\}$ be set such that $\boldsymbol{c}_i \in C_i(\varepsilon_0)$ for $i \in [d]$ and such that $\boldsymbol{b}(\varepsilon_0) \in \operatorname{pos}(C)$. Then, the set $C' = \{\boldsymbol{p} \mid i \in [d],\, \boldsymbol{c}_i \in P_{\varepsilon_0}(\boldsymbol{p})\}$ ray-embraces $\boldsymbol{b}$.*

Proof. We prove the statement by letting ε go continuously from ε_0 to 0. This corresponds to moving the points in C and $\boldsymbol{b}(\varepsilon)$ continuously from their perturbed positions back to their original positions. We argue that throughout this motion, $\boldsymbol{b}(\varepsilon)$ cannot escape the embrace of the colorful choice.

The coordinates of the points in C are defined by polynomials in the parameter ε, and we write $C(\varepsilon)$ for the parametrized points. Then, $C = C(\varepsilon_0)$ and $C' = C(0)$. By Lemma B.3, for all $\varepsilon \in (0, \varepsilon_0]$, the point $\boldsymbol{b}(\varepsilon)$ does not lie in any linear subspace spanned by $d-1$ points from $C(\varepsilon)$. It follows that initially $\boldsymbol{b}(\varepsilon_0) \in \operatorname{int}\operatorname{pos}(C(\varepsilon_0))$ and therefore $\boldsymbol{b}(\varepsilon) \in \operatorname{int}\operatorname{pos}(C(\varepsilon))$ for all $\varepsilon \in (0, \varepsilon_0]$. Assume now that $\boldsymbol{b}(0) \notin \operatorname{pos}(C(0))$. Then, there exists a hyperplane h through $\boldsymbol{0}$ that strictly separates $\boldsymbol{b}(0)$ from $C(0)$. Because the ℓ_2-distance between h and any point in $C(0) \cup \{\boldsymbol{b}(0)\}$ is positive, there is a $\tau \in (0, \varepsilon_0)$ such that h separates $\boldsymbol{b}(\tau)$ from $C(\tau)$, and hence also from $\operatorname{pos}(C(\tau))$. This is impossible, since we showed that $\boldsymbol{b}(\varepsilon) \in \operatorname{int}\operatorname{pos}(C(\varepsilon))$ for all $\varepsilon \in (0, \varepsilon_0]$. □

We can now combine the previous lemmas to obtain our desired result on equivalent instances for ColorfulCarathéodory.

Lemma B.6. *Let $I = (C_1', \ldots, C_d', \boldsymbol{b}')$ be an instance of* ColorfulCarathéodory*, where $C_i' \subset \mathbb{Q}^d$ ray-embraces the point $\boldsymbol{b}' \in \mathbb{Q}^d$ for all $i \in [d]$. Then, we can construct in polynomial time an instance $I^{\approx} = (C_1^{\approx}, \ldots, C_d^{\approx}, \boldsymbol{b}^{\approx})$ of* ColorfulCarathéodory *with properties (P1)–(P3).*

Proof. We construct the point sets $C_1(\varepsilon_0), \ldots, C_d(\varepsilon_0)$ and the point $\boldsymbol{b}(\varepsilon_0)$ as discussed above. Since $\log \varepsilon_0^{-1}$ is polynomial in the size of I, this needs polynomial time. By Lemma B.4, each color class $C_i(\varepsilon_0)$ ray-embraces $\boldsymbol{b}(\varepsilon_0)$, so we can apply Carathéodory's theorem to reduce the size of $C_i(\varepsilon_0)$ to d while maintaining the property that $\boldsymbol{b}(\varepsilon_0)$ is ray-embraced. Again, we need only polynomial time for this step. Finally, as described at the beginning of this section, we rescale the points to lie on the integer grid in polynomial time. Let $C_i^{\approx}$ denote the resulting point set for $C_i(\varepsilon_0)$, where $i \in [d]$, and let $\boldsymbol{b}^{\approx}$ be the point $\boldsymbol{b}(\varepsilon_0)$ scaled to the integer grid. Then, properties (P1)–(P3) are direct consequences of this construction and Lemmas B.3, B.4, and B.5. ☐

C The Colorful Carathéodory Theorem is in PLS

C.1 The Complexity Class PLS

The complexity class *polynomial-time local search* (PLS) [12, 1, 21] captures the complexity of local-search problems that can be solved by a local-improvement algorithm, where each improvement step can be carried out in polynomial time, however the number of necessary improvement steps until a local optimum is reached may be exponential. The existence of a local optimum is guaranteed as the progress of the algorithm can be measured using a potential function that strictly decreases with each improvement step.

More formally, a problem in PLS is a relation $\mathcal{R}$ between a set of *problem instances* $\mathcal{I} \subseteq \{0,1\}^\star$ and a set of *candidate solutions* $\mathcal{S} \subseteq \{0,1\}^\star$. Assume further the following.

- The set $\mathcal{I}$ is polynomial-time verifiable. Furthermore, there exists an algorithm that, given an instance $I \in \mathcal{I}$ and a candidate solution $s \in \mathcal{S}$, decides in time poly($|I|$) whether s is a *valid* candidate solution for I. In the following, we denote with $\mathcal{S}_\mathcal{I} \subseteq \mathcal{S}$ the set of valid candidate solutions for a fixed instance I.
- There exists a polynomial-time algorithm that on input $I \in \mathcal{I}$ returns a valid candidate solution $s \in \mathcal{S}_\mathcal{I}$. We call s the *standard solution*.
- There exists a polynomial-time algorithm that on input $I \in \mathcal{I}$ and $s \in \mathcal{S}_\mathcal{I}$ returns a set $N_{I,s} \subseteq \mathcal{S}_\mathcal{I}$ of valid candidate solutions for I. We call $N_{I,s}$ the *neighborhood* of s.
- There exists a polynomial-time algorithm that on input $I \in \mathcal{I}$ and $s \in \mathcal{S}_I$ returns a number $c_{I,s} \in \mathbb{Q}$. We call $c_{I,s}$ the *cost* of s.

We say a candidate solution $s \in \mathcal{S}$ is a *local optimum* for an instance $I \in \mathcal{I}$ if $s \in \mathcal{S}_I$ and for all $s' \in N_{I,s}$, we have $c_{I,s} \leq c_{I,s'}$ in case of a minimization problem, and $c_{I,s} \geq c_{I,s'}$ in case of a maximization problem. The relation $\mathcal{R}$ then consists of all pairs (I, s) such that s is a local optimum for I. This formulation implies a simple algorithm, that we call the *standard algorithm*: begin with the standard solution, and then repeatedly invoke the neighborhood-algorithm to improve the current solution until this is not possible anymore. Although each iteration of this algorithm can be carried out in polynomial time, the total number of iterations may be exponential. There are straightforward examples in which this algorithm takes exponential time and even more, there are PLS-problems for which it is PSPACE-complete to compute the solution that is returned by the standard algorithm [1, Lemma 15].

Similar to PPAD, each problem instance I of a PLS-problem can be seen as a simple graph searching problem on a graph $G_I = (V, E)$. The set of nodes is the set of valid candidate solutions for I and there is a directed edge from $u \in \mathcal{S}_I$ to $v \in \mathcal{S}_I$ if $v \in N_{I,u}$ and $c_{I,v} < c_{I,u}$ if it is a minimization problem, and otherwise if $c_{I,v} > c_{I,u}$. Then, the set of local optima for I is precisely the set of sinks in G_I. Because the costs induce a topological ordering of the graph, at least one sinks exists.

C.2 A PLS Formulation of the Colorful Carathéodory Problem

The proof of the colorful Carathéodory theorem by Bárány [4] admits a straightforward formulation of ColorfulCarathéodory as a PLS-problem. The only difficulty resides in the computation of the potential function: given a set of d points $C \subset \mathbb{Q}^d$ and a point $\boldsymbol{b} \in \mathbb{Q}^d$, we need to be able to compute the point $\boldsymbol{p}^\star \in \operatorname{pos}(C)$ with minimum ℓ_2-distance to $\boldsymbol{b}$ in polynomial time. This problem can be solved with convex quadratic programming.

We say a matrix $B \in \mathbb{R}^{n\times n}$ is *positive semidefinite* if B is symmetric and for all $\boldsymbol{x} \in \mathbb{R}^n$, we have $\boldsymbol{x}^T B\boldsymbol{x} \geq 0$. Then, a *convex quadratic program* is given by

$$\begin{aligned} Q: \quad & \min c(\boldsymbol{x}) \\ & \text{s.t. } A\boldsymbol{x} = \boldsymbol{b}, \\ & \boldsymbol{x} \geq \boldsymbol{0}, \end{aligned}$$

where $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{b} \in \mathbb{Q}^d$, $A \in \mathbb{Q}^{d\times n}$, and the cost function $c : \mathbb{R}^n \mapsto \mathbb{R}$ is defined as

$$c(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^T B\boldsymbol{x} + \boldsymbol{q}^T\boldsymbol{x},$$

where the matrix $B \in \mathbb{Q}^{n\times n}$ is positive semidefinite and $\boldsymbol{q} \in \mathbb{Q}^n$. We say a vector $\boldsymbol{x} \in \mathbb{R}^n$ is a *feasible solution* for Q if $A\boldsymbol{x} = \boldsymbol{b}$ and $\boldsymbol{x} \geq \boldsymbol{0}$. Furthermore, we say feasible solution $\boldsymbol{x} \in \mathbb{R}^n$ is *optimal* for Q if there is no feasible solution $\boldsymbol{x'} \in \mathbb{R}^n$ such that $c(\boldsymbol{x'}) < c(\boldsymbol{x})$. Convex quadratic programs are known to be solvable in $O(\operatorname{poly}(d, n)L)$ time, where L is the length of the quadratic program in binary [15, 14].

Lemma C.1. *Let $C \subset \mathbb{Q}^d$ be a set of size d and let $\boldsymbol{b} \in \mathbb{Q}^d$ be a point such that C and $\boldsymbol{b}$ can be encoded with L bits. Then, we can compute the point $\boldsymbol{p}^\star \in \operatorname{pos}(C)$ with minimum ℓ_2-distance to $\boldsymbol{b}$ in $O(\operatorname{poly}(d)L)$ time.*

Proof. First, we observe that it is sufficient to compute the point $\boldsymbol{p}^\star \in \operatorname{pos}(C)$ such that

$$\|\boldsymbol{p}^\star - \boldsymbol{b}\|_2^2 = \sum_{i=1}^{d}(\boldsymbol{p}^\star - \boldsymbol{b})_i^2$$

is minimum. Let A be the matrix

$$A = \left(\begin{array}{cccccccc|c|c} \boxed{1 \;\; -1} & & & & & & & 0 & & \\ & & \boxed{1 \;\; -1} & & & & & & C & -\boldsymbol{b} \\ & & & & \ddots & & & & & \\ 0 & & & & & & \boxed{1 \;\; -1} & & & \\ \hline 0 & & & \cdots & & & & & 0 & 1 \end{array}\right) \in \mathbb{Q}^{(d+1)\times(3d+1)}$$

and let $\boldsymbol{b'}$ denote the vector

$$\boldsymbol{b'} = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \hline 1 \end{pmatrix} \in \mathbb{Q}^{d+1}.$$

Furthermore, let $\boldsymbol{x} \in \mathbb{R}^{3d+1}$ be a feasible solution to the linear system

$$A\boldsymbol{x} = \boldsymbol{b'},\ \boldsymbol{x} \geq \boldsymbol{0} \tag{8}$$

and let $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_d$ denote the points in C ordered according to their respective column indices in A. Write $\boldsymbol{x}$ as

$$\boldsymbol{x} = \begin{pmatrix} x_1^+ & x_1^- & x_2^+ & x_2^- & \ldots & x_d^+ & x_d^- & | & \psi_1 & \ldots & \psi_d & | & x_{\boldsymbol{b}} \end{pmatrix}^T \in \mathbb{R}^{3d+1},$$

where $x_i^+, x_i^- \in \mathbb{R}_+$ for $i \in [d]$, $\psi_i \in \mathbb{R}_+$ for $i \in [d]$, and $x_{\boldsymbol{b}} \in \mathbb{R}_+$. Since $\boldsymbol{x} \geq 0$, the point

$$\boldsymbol{p} = \sum_{i=1}^{d} \psi_i \boldsymbol{c}_i$$

is contained in the positive span of C. Furthermore, by the last equality of (8), we have $x_{\boldsymbol{b}} = 1$ and thus for $i \in [d]$, the ith equality of (8) is equivalent to

$$x_i^+ - x_i^- = (\boldsymbol{p})_i - (\boldsymbol{b})_i. \tag{9}$$

Now, let B' denote the matrix

$$B' = \begin{pmatrix} \boxed{\begin{matrix} 1 & -1 \\ -1 & 1 \end{matrix}} & & & 0 \\ & \boxed{\begin{matrix} 1 & -1 \\ -1 & 1 \end{matrix}} & & \\ & & \ddots & \\ 0 & & & \boxed{\begin{matrix} 1 & -1 \\ -1 & 1 \end{matrix}} \end{pmatrix} \in \mathbb{Q}^{(2d)\times(2d)}$$

and set

$$B = \left(\begin{array}{c|c} 2B' & Z_{(2d)\times(d+1)} \\ \hline \multicolumn{2}{c}{Z_{(d+1)\times(3d+1)}} \end{array} \right) \in \mathbb{Q}^{(3d+1)\times(3d+1)},$$

where $Z_{a\times b} \in \mathbb{Q}^{a\times b}$ denotes the all-0 matrix with a rows and b columns. We claim that $\frac{1}{2}\boldsymbol{x}^T B\boldsymbol{x} = \|\boldsymbol{p} - \boldsymbol{b}\|_2^2$. Indeed, by definition of B and using (9), we have

$$\frac{1}{2}\boldsymbol{x}^T B\boldsymbol{x} = \sum_{i=1}^{d} x_i^+ \left(x_i^+ - x_i^-\right) + x_i^- \left(x_i^- - x_i^+\right) = \sum_{i=1}^{d} \left(x_i^+ - x_i^-\right)^2 = \sum_{i=1}^{d} ((\boldsymbol{p})_i - (\boldsymbol{b})_i)^2 = \|\boldsymbol{p} - \boldsymbol{b}\|_2^2.$$

Because B is symmetric, this further implies that B is positive semidefinite.

Let now $\boldsymbol{x}^\star$ be an optimal solution to the convex quadratic program

$$\begin{aligned} \min\ & \frac{1}{2}\boldsymbol{x}^T B\boldsymbol{x} \\ \text{s.t. } & A\boldsymbol{x} = \boldsymbol{b}, \\ & \boldsymbol{x} \geq \boldsymbol{0}. \end{aligned}$$

Then, the point

$$\boldsymbol{p}^\star = \sum_{i=2d+1}^{3d} (\boldsymbol{x}^\star)_i\, \boldsymbol{c}_i \in \mathbb{Q}^d$$

is contained in the positive span of C. Moreover, since $\frac{1}{2}(\boldsymbol{x}^\star)^T B\boldsymbol{x}^\star = \|\boldsymbol{p}^\star - \boldsymbol{b}\|_2^2$ is minimum over all feasible solutions and hence over all points in the positive span of C, $\boldsymbol{p}^\star$ is the point in $\mathrm{pos}(C)$ with minimum ℓ_2-distance to $\boldsymbol{b}$. Using the algorithm from [14] or [15], we can compute $\boldsymbol{p}^\star$ in $O(\mathrm{poly}(d)L)$ time. ☐

Having an algorithm to compute the potential function in polynomial time, we only need to translate the above proof of the colorful Carathéodory theorem to the language of PLS.

Theorem C.2. *The problems* ColorfulCarathéodory*,* Centerpoint*,* Tverberg*, and* SimplicialCenter *are in PPAD ∩ PLS.*

Proof. By Theorem 4.9, ColorfulCarathéodory is in PPAD. We now give a formulation of ColorfulCarathéodory as a PLS-problem. Then statement is then implied by Lemma A.7.

The set of problem instances $\mathcal{I}$ consists of all tuples $(C_1, \dots, C_d, \boldsymbol{b})$, where $d \in \mathbb{N}$, $\boldsymbol{b} \in \mathbb{Q}^d$, $\boldsymbol{b} \neq \boldsymbol{0}$, and for all $i \in [d]$, we have $C_i \subset \mathbb{Q}^d$ and C_i ray-embraces $\boldsymbol{b}$. The set of candidate solutions $\mathcal{S}$ then consists of all d-sets $C \subset \mathbb{Q}^d$, where $d \in \mathbb{N}$. Furthermore, for a given instance $I = (C_1, \dots, C_d, \boldsymbol{b})$, we define the set of valid candidate solutions $\mathcal{S}_I$ as the set of all colorful choices with respect to $C_1, \dots, C_d$. Using linear programming, we can check whether a given tuple $I = (C_1, \dots, C_d, \boldsymbol{b})$ is contained in $\mathcal{I}$ and clearly, we can check in polynomial time whether a set $C \subset \mathbb{Q}^d$ is a colorful choice with respect to I and hence whether $C \in \mathcal{S}_I$.

Let now $I \in \mathcal{I}$ be a fixed instance and $s \in \mathcal{S}_I$ a valid candidate solution. We then define the neighborhood $N_{I,s}$ of s as the set of all colorful choices that can be obtained by swapping one point in s with another point of the same color. The set $N_{I,s}$ can be constructed in polynomial time.

We define the cost $c_{I,s}$ of a colorful choice s as the minimum ℓ_2-distance of a point in $\operatorname{pos}(s)$ to $\boldsymbol{b}$. Using the algorithm from Lemma C.1, we can compute $c_{I,s}$ in polynomial time. Finally, we set the standard solution the colorful choice that consists of the first point from each color class. □

D The Polytopal Complex

We begin with the following standard lemma that bounds the bit-complexity of basic feasible solutions for a linear program.

Lemma D.1. *Let $L : A\boldsymbol{x} = \boldsymbol{b}$ be a linear system, where $A \in \mathbb{Z}^{d \times n}$ and $\boldsymbol{b} \in \mathbb{Z}^d$. Furthermore, let B be a feasible basis for L and let $\boldsymbol{x}$ be the corresponding basic feasible solution. Let m denote the largest absolute value of the entries in A and $\boldsymbol{b}$, and set $N = d!m^d$. Then for $i \in \operatorname{ind}(B)$, we have $|(\boldsymbol{x})_i| = \frac{n_i}{|\det A_{\operatorname{ind}(B)}|}$, where $n_i \in [N]_0$ and $\left|\det A_{\operatorname{ind}(B)}\right| \in [N]$. For $i \in [n] \setminus \operatorname{ind}(B)$, we have $(\boldsymbol{x})_i = 0$.*

Proof. Set $A' = A_{\operatorname{ind}(B)}$. By definition of a feasible basis, we have $\det A' \neq 0$, and by definition of a basic feasible solution $\boldsymbol{x}$, we have $A'\boldsymbol{x}_{\operatorname{ind}(B)} = \boldsymbol{b}$ with $\boldsymbol{x} \geq \boldsymbol{0}$ and $(\boldsymbol{x})_j = 0$ for $j \in [n] \setminus \operatorname{ind}(B)$. Applying Cramer's rule [18], we can express the ith coordinate of $\boldsymbol{x}_{\operatorname{ind}(B)}$ as $\det A'_i / \det A'$, where $i \in [d]$ and A'_i is the matrix that we obtain by replacing the ith column of A' with $\boldsymbol{b}$. Using the Leibniz formula, we can bound the determinant:

$$|\det A'| = \left| \sum_{\sigma \in S_d} \operatorname{sgn}(\sigma) \prod_{i=1}^{d} (A')_{i,\sigma(i)} \right| \leq d!\, m^d = N.$$

And similarly, $|\det A'_i| \leq N$ can be obtained. Because $\boldsymbol{x}$ is a basic feasible solution, we have

$$\frac{\det A'_i}{\det A'} = (\boldsymbol{x})_i \geq 0.$$

Moreover, since A' and $\boldsymbol{b}$ contain only integer entries, the determinants $\det A'$ and $\det A'_i$ are integers. The implies the statement. □

Next, using the techniques from Section B, we can show that a deterministic perturbation of polynomial bit-complexity ensures a non-degenerate intersection of the parameter regions with $\mathcal{M}$.

Lemma D.2. *There exists a constant $c \in \mathbb{N}$ with $c \geq 3$ such that for $\varepsilon = N^{-cd}$ the following holds. Let B be an arbitrary but fixed feasible basis of L^{CC}. Let $h_j \subset \mathbb{R}^d$ denote the hyperplane*

$$h_j = \left\{ \boldsymbol{\mu} \in \mathbb{R}^d \,\middle|\, \left(\hat{\boldsymbol{r}}_{B,\boldsymbol{c}_{\boldsymbol{\mu}}}\right)_j = 0 \right\},$$

and set $H_\Phi = \{h_j \mid j \in [d^2] \setminus \operatorname{ind}(B)\}$. Furthermore, let $H_\Box$ denote the set of supporting hyperplanes for the facets of the unit cube in $\mathbb{R}^d$. Then, for all k-subsets H' of $H_\Phi \cup H_\Box$, the intersection $\bigcap_{h \in H'} h$ is either empty or has dimension $d-k$. In particular, if $k > d$, the intersection must be empty.

Proof. Let H' be a k-subset of $H_\Phi \cup H_\Box$, and suppose that $\bigcap_{h\in H'} h \neq \emptyset$. We denote with $H'_\Phi = H' \cap H_\Phi$ the hyperplanes from H_Φ and similarly, we denote with $H'_\Box = H' \cap H_\Box$ the hyperplanes from $H_\Box$. Set $R = [d^2] \setminus \operatorname{ind}(B)$ and let $\phi_1 < \cdots < \phi_n \in R$ be the indices such that $H'_\Phi = \{h_{\phi_1}, \ldots, h_{\phi_n}\}$, where $n = |H'_\Phi|$. Then the intersection $\bigcap_{i=1}^n h_{\phi_i}$ is the solution space to the system of linear equations

$$\left((\boldsymbol{c}_{\boldsymbol{\mu}})_R - \left(A^{-1}_{\operatorname{ind}(B)} A_R\right)^T (\boldsymbol{c}_{\boldsymbol{\mu}})_{\operatorname{ind}(B)} \right)_{\operatorname{rank}_R(\phi_i)} = 0 \text{ for } i \in [n], \tag{10}$$

where $\operatorname{rank}_R(\phi_i)$ denotes the rank of ϕ_i in R. We write $\operatorname{ind}(B) = \{\beta_1, \ldots, \beta_d\}$, with $\beta_1 < \cdots < \beta_d$ and $\boldsymbol{a}_i = \left(A^{-1}_{\operatorname{ind}(B)} A_R\right)_{\operatorname{rank}_R(\phi_i)}$, for $i \in [n]$. Then, (10) is equivalent to

$$-dN^2 (\boldsymbol{\mu})_{\operatorname{col}(\phi_i)} + dN^2 \boldsymbol{a}_i^T \begin{pmatrix} (\boldsymbol{\mu})_{\operatorname{col}(\beta_1)} \\ \vdots \\ (\boldsymbol{\mu})_{\operatorname{col}(\beta_d)} \end{pmatrix} = -1 - dN^2 - \varepsilon^{\phi_i} + \boldsymbol{a}_i^T \begin{pmatrix} 1 + dN^2 + \varepsilon^{\beta_1} \\ \vdots \\ 1 + dN^2 + \varepsilon^{\beta_d} \end{pmatrix} \text{ for } i \in [n], \tag{11}$$

where $\operatorname{col}(\phi_i)$ and $\operatorname{col}(\beta_i)$ denote the colors of the columns with indices ϕ_i and β_i, respectively. Thus, (11) is of the form

$$A_\Phi \boldsymbol{\mu} = \boldsymbol{b}_\Phi, \tag{12}$$

where $A_\Phi \in \mathbb{Q}^{n \times d}$ and the polynomials $(\boldsymbol{b}_\Phi)_i$, $i \in [n]$, are $(\phi_1, \phi_2, \ldots, \phi_n)$-separated with gap 0. The entries of A_Φ are not necessarily integers due to the occurrence of $A^{-1}_{\operatorname{ind}(B)}$ in the vectors $\boldsymbol{a}_i$. By Lemma D.1, the fractions in $A^{-1}_{\operatorname{ind}(B)}$ all have the same denominator: $\det A_{\operatorname{ind}(B)} \in \mathbb{Z}$. We set $A'_\Phi = \left(\det A_{\operatorname{ind}(B)}\right) A_\Phi$ and $\boldsymbol{b}'_\Phi = \left(\det A_{\operatorname{ind}(B)}\right) \boldsymbol{b}_\Phi$. Then, the linear system

$$A'_\Phi \boldsymbol{\mu} = \boldsymbol{b}'_\Phi \tag{13}$$

is equivalent to (12), where $A'_\Phi \in \mathbb{Z}^{n \times d}$ and $\left(\boldsymbol{b}'_\Phi\right)_i$ is a polynomial in ε with integer coefficients and a nontrivial monomial of degree ϕ_i for $i \in [n]$. Let m' denote the maximum absolute value of the coefficients of ε-polynomials in A'_Φ and $\boldsymbol{b}'_\Phi$. Since the absolute value of the entries of A_R is at most N and since by Lemma D.1 the absolute value of the entries in $A^{-1}_{\operatorname{ind}(B)}$ is at most N, there exists a constant $c' \in \mathbb{N}$ such that $m' \leq N^{c'}$ and c' is independent of the choice of B.

Set $n' = |H'_\Box|$. Since we assume that the hyperplanes in H' have a point in common and since $H'_\Box \subseteq H'$, the hyperplanes in $H'_\Box$ fix the values of exactly n' coordinates $(\boldsymbol{\mu})_j$ to either 0 or 1. Let J be the indices of the fixed coordinates and let $J_i \subseteq J$ be the indices of the $(\boldsymbol{\mu})_j$ that are set to i for $i = 0, 1$. Combining this with (13), we can express the intersection of hyperplanes in H' as

$$\left(A'_\Phi\right)_{[d]\setminus J} (\boldsymbol{\mu})_{[d]\setminus J} = \boldsymbol{b}'_\Phi - \sum_{j \in J_1} (A'_\Phi)_j. \tag{14}$$

The matrix $(A'_\Phi)_J$ is an $n \times (d-n')$ integer matrix, whose entries have absolute value at most $N^{c'}$ and the polynomials $p_i = (\boldsymbol{b}'_\Phi - \sum_{j \in J_1} (A'_\Phi)_j)_i$, $i \in [n]$, are $(\phi_1, \phi_2, \ldots, \phi_n)$-separated with gap 0. Then, Lemma B.2 implies that for all $\varepsilon \in \left(0, \frac{1}{2M}\right)$, the right hand vector of (14) cannot lie in the span of $n-1$ columns of the left hand matrix, where $M = d!(d^2+1)\left(N^{c'}\right)^d$. Thus, for $c = \max(3, 2c')$, we have $N^{-cd} \in \left(0, \frac{1}{2M}\right)$. Since we know that (14) has a solution, it follows that the rank of (14) must be n and thus the intersection $\bigcap_{h \in H'} h$ has dimension $d - n - n' = d - k$. □

Note that since c is a constant, the number of bits needed to represent ε is polynomial in the size of the ColorfulCarathéodory instance. We continue by showing that the elements from Q are indeed polytopes and by characterizing precisely their dimension and their facets.

Lemma D.3. *Let $q = \Phi(f) \cap g \neq \emptyset$ be an element from $\mathcal{Q}$, where $f \in \mathcal{F}$ and g is a face of $\mathcal{M}$. Then, q is a simple polytope of dimension* $\dim g - \dim f$*. Moreover, if* $\dim q > 0$*, the set of facets of q can be written as*

$$\left\{\Phi\left(f\right) \cap \breve{g} \neq \emptyset \,\middle|\, \breve{g} \text{ is a facet of } g\right\} \cup \left\{\Phi\left(\hat{f}\right) \cap g \neq \emptyset \,\middle|\, f \text{ is a facet of } \hat{f} \in \mathcal{F}\right\}.$$

Proof. Let B be a feasible basis for a vertex of f. As discussed above, the solution space to the linear system $L^\Phi_{B,f}$ is $\Phi(f)$. We denote with $H^=_{\Phi(f)}$ the set of hyperplanes that are given by the equality constraints

$$(\hat{\boldsymbol{r}}_{B,\boldsymbol{\mu}})_j = 0, \text{ for } j \in \operatorname{supp}(f) \setminus \operatorname{ind}(B),$$

and we denote with $H^-_{\Phi(f)}$ the set of halfspaces that are given by the $d^2 - (d + \dim f)$ inequalities

$$(\hat{\boldsymbol{r}}_{B,\boldsymbol{\mu}})_j \leq 0, \text{ for } j \in \left[d^2\right] \setminus \operatorname{supp}(f)$$

in $L^\Phi_{B,f}$.

Because g is a face of $\mathcal{M}$ and hence of the unit cube, we can write it as the intersection of a set $H^=_g$ of $d - \dim g$ hyperplanes and a set of halfspaces H^-_g, where $H^=_g$ and the boundary hyperplanes from the halfspaces in H^-_g are supporting hyperplanes of facets of the unit cube.

We set $H^= = H^=_g \cup H^=_{\Phi(f)}$ and $H^- = H^-_g \cup H^-_{\Phi(f)}$. Now, q is the intersection of the affine space $S^= = \bigcap_{h \in H^=} h$ with the polyhedron $S^- = \bigcap_{h^- \in H^-} h^-$. Hence, q is a polyhedron and moreover, as $q \subseteq \mathcal{M}$, it is a polytope. By Lemma D.2, the hyperplanes in $H^=$ and the boundary hyperplanes of H^- are in general position, so q is simple.

We now prove $\dim q = \dim g - \dim f$. Because $|H^=_g| = d - \dim g$, $|H^=_{\Phi(f)}| = \dim f$, and by Lemma D.2, we have $H^=_g \cap H^=_{\Phi(f)} = \emptyset$, the set $H^=$ contains $d - \dim g + \dim f$ hyperplanes. Again by Lemma D.2, the hyperplanes from $H^=$ are in general position, and therefore $\dim S^= = \max(\dim g - \dim f, -1)$, where we set $\dim \emptyset = -1$. Since we assume that $q \neq \emptyset$, it follows that $\dim S^= \geq 0$, so in particular $\dim f \leq \dim g$. We show that the dimension does not decrease by intersecting $S^=$ with the halfspaces in H^-. Fix an arbitrary ordering $h^-_1, \ldots, h^-_m$, $m = |H^-|$, of the halfspaces in H^-. For $j = 0, 1, \ldots, m$, let Ψ_j denote the polyhedron that we obtain by intersecting $S^=$ with the first j halfspaces $h^-_1, \ldots, h^-_j$ from H^-. In particular, we have $\Psi_0 = S^=$ and $\Psi_m = q$. Assume for the sake of contradiction that $\dim q < \dim S^=$, and let $j^\star$ be such that $\dim \Psi_{j^\star - 1} = \dim S^=$ and $\dim \Psi_{j^\star} = d_{j^\star} < \dim S^=$. There are three possibilities: (i) $\Psi_{j^\star-1} \cap h^-_{j^\star} = \emptyset$; (ii) $h^-_{j^\star}$ intersects the relative interior of $\Psi_{j^\star-1}$; or (iii) $h^-_{j^\star}$ intersects only the boundary of $\Psi_{j^\star-1}$. Now, since $q \neq \emptyset$, Case (i) is impossible. Since by our assumption, $d_{j^\star} < \dim \Psi_{j^\star-1}$, Case (ii) also cannot occur. Hence, $\Psi_{j^\star}$ is a proper face of $\Psi_{j^\star-1}$. Then, $\Psi_{j^\star}$ is contained in the intersection of the $d - \dim g + \dim f$ hyperplanes from $H^=$ with at least $\dim S^= - d_{j^\star} = \dim g - \dim f - d_{j^\star}$ boundary hyperplanes of $h_1^-, \dots, h_{j^\star-1}^-$, and with the boundary hyperplane of $h_{j^\star}^-$. Thus, the $d_{j^\star}$-dimensional polyhedron $\Psi_{j^\star}$ lies in the intersection of at least $d - d_{j^\star} + 1$ hyperplanes from $H^=$ and bounding hyperplanes from H^-. Hence, the hyperplanes from $H^=$ together with the bounding hyperplanes from H^- are not in general position, a contradiction to Lemma D.2.

We now prove the second part of the statement. Let $\check{q}$ be a facet of q. Since $\dim q > 0$, the facet $\check{q}$ is nontrivial. Then, $\check{q}$ is the intersection of q with a hyperplane $h^\star$ that is a boundary hyperplane of some halfspace in H^-. Let h^- be the halfspace that generates $h^\star$. If $h^- \in H_g^-$, then $\check{g} = g \cap h$ is a facet of g and we have $\check{q} = \Phi(f) \cap \check{g}$. Assume now $h^- \in H_{\Phi(f)}^-$ and let h be defined by the equation $(\hat{\boldsymbol{r}}_{B,\boldsymbol{c}_{\boldsymbol{\mu}}})_j = 0$ for some $j \in \operatorname{supp}(f) \setminus \operatorname{ind}(B)$. Let $\hat{f} \subseteq \mathcal{P}^{\mathrm{CC}}$ be the face that is defined by the columns from A with indices $\operatorname{supp}(f) \cup \{j\}$, and note that f is a facet of $\hat{f}$. Then, we can write $\check{q}$ as

$$\check{q} = h^\star \cap q = h^\star \cap \left(\bigcap_{h \in H_{\Phi(f)}^=} h \cap \bigcap_{h^- \in H_{\Phi(f)}^-} h^- \cap g \right) = \left(h^\star \cap \bigcap_{h \in H_{\Phi(f)}^=} h \cap \bigcap_{h^- \in H_{\Phi(f)}^-} h^- \right) \cap g$$

and thus $\check{q}$ contains all parameter vectors in g for which $\hat{f}$ is optimal.

Now, let $\check{g}$ be a facet of g with $\check{q} = \Phi(f) \cap \check{g} \neq \emptyset$. Then, there exists a boundary hyperplane $h^\star$ from a halfspace in H_g^- such that $\check{q} = h^\star \cap (\bigcap_{h \in H^=} h) \cap (\bigcap_{h^- \in H^-} h^-)$. Clearly, $\check{q}$ is a face of q. Furthermore, since $\check{q} \neq \emptyset$ the first part of the lemma implies

$$\dim \check{q} = \dim \check{g} - \dim f = (\dim g - 1) - \dim f = \dim q - 1 \geq 0.$$

Hence $\check{q}$ is a facet of q. Let now $\hat{f} \in \mathcal{F}$ be a face that has f as a facet with $\check{q} = \Phi(\hat{f}) \cap g \neq \emptyset$. Then there exists a boundary hyperplane $h^\star$ of a halfspace in $H_{\Phi(f)}^-$ such that $\check{q} = h^\star \cap (\bigcap_{h \in H^=} h) \cap (\bigcap_{h^- \in H^-} h^-)$. As before, $\check{q}$ is a face of q and since $\check{q} \neq \emptyset$, we get

$$\dim \check{q} = \dim g - \dim \hat{f} = \dim g - (\dim f + 1) = \dim q - 1 \geq 0.$$

Thus, $\check{q}$ is a facet of q. □

In particular, Lemma D.3 implies that within each k-face of $\mathcal{M}$, the set of parameter vectors that are optimal for some vertex $v \in \mathcal{F}$ is either empty or a k-dimensional polytope and the set of parameter vectors that are optimal for a k-face $f \in \mathcal{F}$ is either empty or a single point. Furthermore, Lemma D.3 immediately bounds the maximum dimensions of faces in $\mathcal{F}$.

The next lemma shows that the intersection of any two polytopes in $\mathcal{Q}$ is again an element in $\mathcal{Q}$.

Lemma D.4. *Let $q_1 = \Phi(f_1) \cap g_1 \in \mathcal{Q}$ and $q_2 = \Phi(f_2) \cap g_2 \in \mathcal{Q}$ be two polytopes with $q_1 \cap q_2 \neq \emptyset$, where f_1, $f_2 \in \mathcal{F}$ and g_1, g_2 are faces of $\mathcal{M}$. Then,*

$$q_1 \cap q_2 = \Phi\left(\hat{f}\right) \cap \check{g},$$

where $\hat{f} \in \mathcal{F}$ is the smallest face of $\mathcal{P}^{\mathrm{CC}}$ that contains f_1 and f_2, and $\check{g} = g_1 \cap g_2$.

Proof. We begin with showing that $\Phi(f_1) \cap \Phi(f_2) = \Phi\left(\hat{f}\right)$. Let $\boldsymbol{\mu} \in \Phi(f_1) \cap \Phi(f_2)$ be a vector. Since $\hat{f}$ is the smallest face of $\mathcal{P}^{\mathrm{CC}}$ that contains f_1 and f_2, the face $\hat{f}$ is optimal for $L_{\boldsymbol{\mu}}^{\mathrm{CC}}$ and thus $\Phi(f_1) \cap \Phi(f_2) \subseteq \Phi\left(\hat{f}\right)$. Let now $\boldsymbol{\mu}$ be a parameter vector from $\Phi\left(\hat{f}\right)$. Since f_1 and f_2 are subfaces of $\hat{f}$, the faces f_1 and f_2 are optimal for $\boldsymbol{\mu}$ and thus we have $\boldsymbol{\mu} \in \Phi(f_1) \cap \Phi(f_2)$. Hence, $\Phi\left(\hat{f}\right) = \Phi(f_1) \cap \Phi(f_2)$. Then, we can express $q_1 \cap q_2$ as

$$q_1 \cap q_2 = (\Phi(f_1) \cap g_1) \cap (\Phi(f_2) \cap g_2) = \Phi\left(\hat{f}\right) \cap \check{g},$$

where $\check{g} = g_1 \cap g_2$. Moreover, since $q_1 \cap q_2 \neq \emptyset$ and $\check{g}$ is a face of $\mathcal{M}$, the face $\hat{f}$ is contained in $\mathcal{F}$. □

Equipped with Lemmas D.3 and D.4, we are now ready to show that $\mathcal{Q}$ is a polytopal complex.

Lemma D.5. *The set $\mathcal{Q}$ is a $(d-1)$-dimensional polytopal complex that decomposes $\mathcal{M}$.*

Proof. Lemma D.3 guarantees that every element $q \in \mathcal{Q}$ is a polytope in $\mathbb{R}^d$ of dimension at most $d-1$. By the second part of Lemma D.3, if $\dim q > 0$, all facets of q and hence inductively all nonempty faces of q are contained in $\mathcal{Q}$. Furthermore, since $\emptyset$ is a face of $\mathcal{M}$, it is contained in $\mathcal{Q}$ as well.

Now, let $q_1, q_2 \in \mathcal{Q}$ be two polytopes. If $q_1 \cap q_2 = \emptyset$, then clearly $q_1 \cap q_2$ is a face of both polytopes q_1 and q_2, so assume $q_1 \cap q_2 \neq \emptyset$. By definition of $\mathcal{Q}$, there are faces $f_1, f_2 \in \mathcal{F}$ and faces g_1, g_2 of $\mathcal{M}$ such that $q_1 = \Phi(f_1) \cap g_1$ and $q_2 = \Phi(f_2) \cap g_2$. Then, we can apply Lemma D.4 to express the intersection of q_1 and q_2 as $\Phi\left(\hat{f}\right) \cap \check{g}$. Since $\hat{f} \in \mathcal{F}$ and since $\check{g}$ is a face of $\mathcal{M}$, $q_1 \cap q_2 \in \mathcal{Q}$. Moreover, as $\hat{f}$ is a superface of f_1 and $\check{g}$ is a face of g_1, a repeated application of Lemma D.3 shows that $q_1 \cap q_2$ is a face of q_1. Similarly, because $\hat{f}$ is a superface of f_2 and $\check{g}$ is a face of g_2, a repeated application of Lemma D.3 proves that $q_1 \cap q_2$ is a face of q_2, as desired. □

A further implication of Lemmas D.3 and D.4 is that each polytope in $\mathcal{Q}$ can be represented uniquely as the intersection of a parameter region of a face of $\mathcal{P}^{\mathrm{CC}}$ and a face of $\mathcal{M}$.

Lemma D.6. *Let $q \in \mathcal{Q}$ be a polytope. Then, there exists unique pair of faces f, g, where $f \in \mathcal{F}$ and g is a face of $\mathcal{M}$, such that $q = \Phi(f) \cap g$.*

Proof. Let f_1, f_2 be two faces of $\mathcal{P}^{\mathrm{CC}}$ and let g_1, g_2 be two faces of $\mathcal{M}$ such that

$$q = \Phi(f_1) \cap g_1 = \Phi(f_2) \cap g_2.$$

Then, by Lemma D.4, we can write q as $\Phi\left(\hat{f}\right) \cap \check{g}$, where $\hat{f} \in \mathcal{F}$ is the smallest face in $\mathcal{P}^{\mathrm{CC}}$ that contains f_1 and f_2 and $\check{g}$ is a face of g_1 and of g_2. If $\hat{f} \neq f_1$ or $\check{g} \neq g_1$, then by Lemma D.3,

$$\dim q = \dim \check{g} - \dim \hat{f} < \dim g_1 - \dim f_1 = \dim q,$$

a contradiction. Hence, we must have $\hat{f} = f_1$ and $\check{g} = g_1$. Similarly, we must have $\hat{f} = f_2$ and $\check{g} = g_2$, and thus $f_1 = f_2$ and $g_1 = g_2$. □

Lemmas 4.2 and 4.3 are now immediate consequences from Lemmas D.4, D.6, and D.4.

We conclude with the proof of Lemma 4.1. For this, we need the following observation that is a direct consequence of Property (P2) of the ColorfulCarathéodory instance.

Observation D.7. *For any feasible basis B of L^{CC}, the coordinates for B in the corresponding basic feasible solution are strictly positive. Equivalently, $\mathcal{P}^{\mathrm{CC}}$ is simple.*

Proof of Lemma 4.1. Let $\boldsymbol{x}^\star$ be the basic feasible solution for $B^\star$ with respect to $L_{\boldsymbol{\mu}}^{\mathrm{CC}}$. For the sake of contradiction, suppose that $B^\star$ contains some vector of $C_{i^\times}$, and let k be the index of the corresponding coordinate in $\boldsymbol{x}^\star$. By Observation D.7 and Lemma D.1, we have $(\boldsymbol{x}^\star)_k \geq 1/N$. Hence,

$$\boldsymbol{c}_{\boldsymbol{\mu}}^T \boldsymbol{x}^\star \geq (\boldsymbol{c}_{\boldsymbol{\mu}})_k (\boldsymbol{x}^\star)_k \geq \left(1 + dN^2\right) (\boldsymbol{x}^\star)_k \geq dN + \frac{1}{N}$$

since $\boldsymbol{c}_{\boldsymbol{\mu}} \geq \mathbf{1}$ and $\boldsymbol{x}^\star \geq \mathbf{0}$. By construction, there is a color $i^\star \in [d]$ such that $(\boldsymbol{c}_{\boldsymbol{\mu}})_j = 1 + \varepsilon^j$ for all columns j with color $i^\star$. Let $\boldsymbol{x}^{(i^\star)}$ be the basic feasible solution for the basis $C_{i^\star}$. By Lemma D.1, $\left(\boldsymbol{x}^{(i^\star)}\right)_j$ is upper bounded by N for all $j \in \operatorname{ind}(C_{i^\star})$, so we can lower bound the costs of $\boldsymbol{x}^{(i^\star)}$ as follows:

$$\boldsymbol{c}_{\boldsymbol{\mu}}^T \boldsymbol{x}^{(i^\star)} = \sum_{j \in \operatorname{ind}(C_{i^\star})} (\boldsymbol{c}_{\boldsymbol{\mu}})_j \left(\boldsymbol{x}^{(i^\star)}\right)_j \leq \sum_{j \in \operatorname{ind}(C_{i^\star})} \left(1 + \frac{1}{N^3}\right) \left(\boldsymbol{x}^{(i^\star)}\right)_j \leq dN + \frac{d}{N^2} < dN + \frac{1}{N},$$

where we use that $0 < \varepsilon \leq N^{-3}$. This contradicts the optimality of $B^\star$. □

E The Barycentric Subdivison – Omitted Proofs

Proof of Lemma 4.4. Let $q_0 \subset \cdots \subset q_{d-1}$ be the chain that corresponds to σ in $\operatorname{sd} \mathcal{Q}_\Delta$. By Lemma 4.2, we can write each polytope $q_i \in \mathcal{Q}_\Delta$ uniquely as $\Phi_\Delta(f_i) \cap g_i$, where $i \in [d-1]_0$, $f_i \in \mathcal{F}$, and $g_i \in \mathcal{S}$. By the definition of the barycentric subdivision and since $\mathcal{Q}_\Delta$ is a $(d-1)$-dimensional polytopal complex, q_{i-1} is a facet of q_i for $i \in [d-1]$. Then, Lemma 4.2 states that either g_{i-1} is a facet of g_i or f_i is a facet of f_{i-1} for $i \in [d-1]$. Because σ is fully-labeled, we must have $f_i \neq f_j$ for all $i, j \in [d-1]_0$ with $i \neq j$. Hence, f_i is a facet of f_{i-1} for $i \in [d-1]$ and thus $g_0 = \cdots = g_{d-1}$. Since $\dim q_{d-1} = d-1$, Lemma 4.2 implies that $\dim f_i = d-1-i$ and hence $|\operatorname{supp}(f_i)| = 2d-1-i$ for $i \in [d-1]_0$. In particular, $\dim f_{d-1} = 0$ and thus the columns from $A_{\operatorname{supp}(f_{d-1})}$ are a feasible basis for L^{CC}. For $i \in [d-1]$, let $a_{i-1} \in \left[d^2\right]$ denote the column index such that $\operatorname{supp}(f_{i-1}) = \operatorname{supp}(f_i) \cup \{a_{i-1}\}$. Since the faces $f_0, \ldots, f_{d-1}$ have pairwise distinct labels and since $|\operatorname{supp}(f_{i-1})| = |\operatorname{supp}(f_i)| + 1$ for $i \in [d-1]$, the column vectors $A_{a_0}, \ldots, A_{a_{d-2}}$ have pairwise distinct colors by the definition of λ (see (5)). Now assume for the sake of contradiction that the columns from $A_{\operatorname{supp}(f_{d-1})}$ are not a colorful feasible basis. Then, there is some color $i^\times \in [d]$ that does not appear in $A_{\operatorname{supp}(f_{d-1})}$ and hence there is some color $i^\star \in [d]$ with $|\operatorname{ind}(C_{i^\star}) \cap \operatorname{supp}(f_{d-1})| \geq 2$. Since there is at most one column with color $i^\times$ among $A_{a_0}, \ldots, A_{a_{d-2}}$, we have $|\operatorname{supp}(f_i) \cap \operatorname{ind}(C_{i^\times})| \leq 1$ for all $i \in [d-1]_0$. Since $\operatorname{supp}(f_i) \supseteq \operatorname{supp}(f_{d-1})$ for $i \in [d-1]_0$ and since $|\operatorname{ind}(C_{i^\star}) \cap \operatorname{supp}(f_{d-1})| \geq 2$, we have $\lambda(f_i) \neq i^\times$ for all $i \in [d-1]_0$, a contradiction to σ being fully-labeled. □

Proof of Lemma 4.5. We begin by showing that the encoding $\operatorname{enc}(\sigma)$ of a simplex $\sigma \in \Sigma_k$ is a valid k-tuple. Let $q_0 \subset \cdots \subset q_{k-1}$ be the corresponding face chain in $\mathcal{Q}_\Delta$ such that the ith vertex of σ is the barycenter of $q_i \in \mathcal{Q}_\Delta$ and $q_i \neq \emptyset$ for $i \in [k-1]_0$. By Lemma 4.2, for each q_i, $i \in [k-1]_0$, there exists a unique pair of faces $f_i \in \mathcal{F}$ and $g_i \in \mathcal{S}$ such that $q_i = \Phi_\Delta(f_i) \cap g_i$. Because $q_{k-1} \neq \emptyset$, we have $\mathcal{M}(q_{k-1}) = \Phi(f_i) \cap g\left(I_0^{(k-1)}, I_1^{(k-1)}\right) \neq \emptyset$. We further observe that $g_i \subset \Delta_{[k]}$. Otherwise we would have $q_i = \Phi_\Delta(f_i) \cap \left(g_i \cap \Delta_{[k]}\right)$ with $g_i \cap \Delta_{[k]} \in \mathcal{S}$, a contradiction to g_i, f_i being the unique pair. Since $q_i \subset \Delta_{[k]}$ for $i \in [k-1]_0$ and since $\dim \Delta_{[k]} = k-1$, we must have $\dim q_i = i$ for $i \in [k-1]$. Then, Lemma 4.2 implies that $\dim g_{k-1} = k-1$ and $\dim f_{k-1} = 0$. In particular, $\operatorname{supp}(f_{k-1})$ is the index set of a feasible basis and $\left|I_0^{(k-1)} \cup I_1^{(k-1)}\right| = d-k+1$. Because $g_{k-1} \subset \Delta_{[k]}$, we have

$[d] \setminus [k] \subseteq I_0^{(k-1)}$ and since g_{k-1} is the projection of a face of $\mathcal{M}$, the set $I_1^{(k-1)}$ is nonempty. Thus, $I_0^{(k-1)} = [d] \setminus [k]$ and $\left|I_1^{(k-1)}\right| = 1$.

Let now $i \in [k-1]$ be a fixed index and write $\operatorname{enc}(q_{i-1}) = \left(\operatorname{supp}(f_{i-1}), I_0^{(i-1)}, I_1^{(i-1)}\right)$ and $\operatorname{enc}(q_i) = \left(\operatorname{supp}(f_i), I_0^{(i)}, I_1^{(i)}\right)$. Since q_{i-1} is a facet of q_i, Lemma 4.2 implies that either (a) f_i is a facet of f_{i-1} and $g_{i-1} = g_i$ or (b) $f_{i-1} = f_i$ and g_{i-1} is a facet of g_i. In Case (a), we have $\operatorname{supp}(f_{i-1}) = \operatorname{supp}(f_i) \cup \{a_{i-1}\}$ and $I_0^{(i-1)} = I_0^{(i)}$ as well as $I_1^{(i-1)} = I_1^{(i)}$, where $a_{i-1} \in [d^2] \setminus \operatorname{supp}(f_i)$. In Case (b), we have $\operatorname{supp}(f_{i-1}) = \operatorname{supp}(f_i)$. Furthermore, since $\mathcal{M}(g_{i-1})$ is a facet of $\mathcal{M}(g_i)$, we either have $I_0^{(i-1)} = I_0^{(i)} \cup \{j_{i-1}\}$ and $I_1^{(i-1)} = I_1^{(i)}$, or $I_1^{(i-1)} = I_1^{(i)} \cup \{j_{i-1}\}$ and $I_0^{(i-1)} = I_0^{(i)}$, for an index $j_{i-1} \in [d] \setminus \left(I_0^{(i)} \cup I_1^{(i)}\right)$. Thus, $\operatorname{enc}(\sigma)$ is a valid k-tuple.

We now show that enc is a bijection. Let $\sigma_1, \sigma_2 \in \Sigma_k$ be two simplices. Since the barycenters of the polytopes in a polytopal complex are pairwise distinct, the face chains in $\mathcal{Q}_\Delta$ that corresponds to σ_1 and σ_2 must differ in at least one face. Then, (6) together with Lemma 4.2 directly implies that $\operatorname{enc}(\sigma_1) \neq \operatorname{enc}(\sigma_2)$.

Let now $T = (Q_0, \ldots, Q_{k-1})$, $k \in [d-1]$, be a valid k-tuple, where $Q_i = \left(S^{(i)}, I_0^{(i)}, I_1^{(i)}\right)$. For $i \in [k-1]_0$, let $g_i' = g\left(I_0^{(i)} \cup I_1^{(i)}\right)$ be the subset of $\mathcal{M}$ that is defined by the index sets $I_0^{(i)}, I_1^{(i)}$. Since $[d] \setminus [k] \subseteq I_0^{(i)}$ for all $i \in [k-1]_0$, the projection $g_i = \Delta(g_i')$ is a subset of $\Delta_{[k]}$. Moreover, since $I_1^{(i)} \neq \emptyset$ for $i \in [k-1]_0$, the set g_i' is a face of $\mathcal{M}$ and hence $g_i \in \mathcal{S}$. Furthermore, since the columns in $A_{S^{(k-1)}}$ are a feasible basis, they define a vertex f_{k-1}. Because $S^{(k-1)} \subseteq S_i$ for $i \in [k-1]_0$, the index set S_i is the support of a face $f_i \in \mathcal{F}$. Set $q_i = \Phi_\Delta(f_i) \cap g_i \in \mathcal{Q}$ for $i \in [k-1]_0$. Because $g_i \subset \Delta_{[k]}$, the polytope q_i is also contained in $\Delta_{[k]}$. By Property (i) of a valid sequence, the intersection $\Phi(f_{k-1}) \cap g_{k-1}'$ is nonempty and hence its projection q_{k-1} onto Δ is nonempty. Then, Lemma 4.2 states that $\dim q_{k-1} = k-1$. Moreover by Lemma 4.2 and properties (ii.a) and (ii.b) of T, either g_{i-1} is a facet of g_i or f_i is a facet of f_{i-1} for $i \in [k-1]$. Thus by Lemma 4.2, q_{i-1} is a facet of q_i, $i \in [k-1]$. Then, $\dim q_i = i$ for all $i \in [k-1]_0$ and hence the face chain $q_0 \subset \cdots \subset q_{k-1}$ defines a $(k-1)$-simplex $\sigma \in \Sigma_k$ with $\operatorname{enc}(\sigma) = T$. □

Proof of Lemma 4.6. Clearly, we can check if T fulfills all syntactic requirements on valid k-tuples in polynomial time. Furthermore, we can check in polynomial time whether the columns B from $A_{S^{(k-1)}}$ are a feasible basis for a vertex f. Finally, we express $\Phi(f) \cap g\left(I_0^{(k-1)}, I_1^{(k-1)}\right)$ as the solution space to the linear system $L_{B,f}^{\mathrm{CC}}$ extended by the constraints $\boldsymbol{\mu} \in g\left(I_0^{(k-1)}, I_1^{(k-1)}\right)$. Then, we can check in polynomial time whether this system has a solution. □

The key for Lemma 4.7 is the following lemma that guarantees that simplices with facets in common have a similar encoding.

Lemma E.1. *Let $\sigma, \sigma' \in \Sigma_k$ be two simplices, where $k \in [d]$. Then, σ and σ' share a facet if and only if the tuples* $\operatorname{enc}(\sigma)$ *and* $\operatorname{enc}(\sigma')$ *agree in all but one position. Furthermore, let $\sigma \in \Sigma_k$ and $\hat{\sigma} \in \Sigma_{k+1}$ be two simplices, where $k \in [d-1]_0$. Write* $\operatorname{enc}(\sigma)$ *as*

$$\operatorname{enc}(\sigma) = \left(Q_0, \ldots, Q_{k-1} = \left(S^{(k-1)}, I_0^{(k-1)}, I_1^{(k-1)}\right)\right).$$

Then, σ is a facet of $\hat{\sigma}$ if and only if

$$\operatorname{enc}(\hat{\sigma}) = \left(Q_0, \ldots, Q_{k-1}, \left(S^{(k-1)}, I_0^{(k-1)} \setminus \{k+1\}, I_1^{(k-1)}\right)\right).$$

Proof. Let $\sigma, \sigma' \in \Sigma_k$ be two simplices and let $q_0 \subset \cdots \subset q_{k-1}$ and $q'_0 \subset \cdots \subset q'_{k-1}$ be the corresponding face chains in $\mathcal{Q}_\Delta$. Then σ and σ' share a facet if and only if the face chains agree on all but one position and hence if and only if $\operatorname{enc}(\sigma)$ and $\operatorname{enc}(\sigma')$ agree on all but one position.

Let now $\sigma \in \Sigma_k$ and $\hat{\sigma} \in \Sigma_{k+1}$ be two simplices. Let $q_0 \subset \cdots \subset q_{k-1}$ be the face chain in $\mathcal{Q}_\Delta$ that corresponds to σ with $\dim q_i = i$ for $i \in [k-1]_0$. Similarly, let $\hat{q}_0 \subset \cdots \subset \hat{q}_k$ be the face chain in $\mathcal{Q}_\Delta$ that corresponds to $\hat{\sigma}$ with $\dim \hat{q}_i = i$ for $i \in [k]_0$. Furthermore, we write $\operatorname{enc}(q_{k-1}) = \left(S^{(k-1)}, I_0^{(k-1)}, I_1^{(k-1)}\right)$ and $\operatorname{enc}(\hat{q}_k) = \left(S^{(k)}, I_0^{(k)}, I_1^{(k)}\right)$. Then, σ is a facet of $\hat{\sigma}$ if and only if the faces $q_0, \ldots, q_{k-1}$ appear in the face chain of $\hat{\sigma}$ and hence if and only if $q_i = q'_i$ for $i \in [k-1]_0$. Moreover, since by Lemma 4.5 the encodings $\operatorname{enc}(\sigma)$ and $\operatorname{enc}(\hat{\sigma})$ are valid tuples, the columns of $A_{S^{(k-1)}}$ and $A_{S^{(k)}}$ are feasible bases. Since $S^{(k-1)} \subseteq S^{(k)}$ by Property (ii) of valid tuples, we must have $S^{(k-1)} = S^{(k)}$. Moreover, by Property (i), we have $I_0^{(k-1)} = [d] \setminus [k]$, $I_0^{(k)} = [d] \setminus [k+1]$, and $\left|I_1^{(k-1)}\right| = \left|I_1^{(k)}\right| = 1$. Because of Property (ii), the index set $I_1^{(k-1)}$ is a subset of $I_1^{(k)}$ and hence $I_1^{(k-1)} = I_1^{(k)}$. We conclude that

$$\operatorname{enc}(\hat{\sigma}) = \left(\operatorname{enc}(q_0), \ldots, \operatorname{enc}(q_{k-1}), \left(S^{(k-1)}, I_0^{(k-1)} \setminus \{k+1\}, I_1^{(k-1)}\right)\right),$$

as claimed. □

Proof of Lemma 4.7. We begin with the first problem. By Lemma E.1, if there is a simplex $\sigma' \in \Sigma_k$ that shares the facet $\operatorname{conv}\{\boldsymbol{v}_j \mid j \in [k-1]_0,\, j \neq i\}$ with σ, the encodings $\operatorname{enc}(\sigma)$ and $\operatorname{enc}(\sigma')$ agree on all but one position. Thus, there are only polynomially many possibilities for the encoding of $\operatorname{enc}(\sigma')$ that we can check in polynomial time with the algorithm from Lemma 4.6. Furthermore, Lemma E.1 directly implies polynomial-time algorithms for the second and third problem. □

F The PPAD Graph

We begin by characterizing by showing that the graph consists only of paths and cycles and by characterizing the degree one nodes.

Proof of Lemma 4.8. Let $\operatorname{enc}(\sigma) \in V_k$ be the encoding of a simplex $\sigma \in \Sigma_k$. If $\sigma \in \Sigma_1$ then $\deg \operatorname{enc}(\sigma) = 1$ since the only adjacent node is the encoding of the simplex in Σ_2 with σ as a facet. Similarly, if $\operatorname{enc}(\sigma) \in V_d$ with $\lambda(\sigma) = [d]$, then $\deg \operatorname{enc}(\sigma) = 1$ since the only adjacent node is either the encoding of the single $[d-1]$-labeled facet of σ or the encoding of the simplex in Σ_d that shares this facet.

If $k > 1$ and σ has two $[k-1]$-labeled facets, then $\deg \operatorname{enc}(\sigma) = 2$ since each $[k-1]$-labeled facet is either shared with another simplex in Σ_k or the facet is itself in Σ_{k-1}. Otherwise, if $k < d$ and $\lambda(\sigma) = [k]$, then we have again $\deg \operatorname{enc}(\sigma) = 2$ as there exists exactly one simplex in Σ_{k+1} with σ as a facet and either the single $[k-1]$-labeled facet of σ is shared with another simplex in Σ_k or it is itself a simplex in Σ_{k-1}. Note that actually Lemma 4.5 implies in this case that the $[k-1]$-labeled facet must be shared with another simplex in Σ_k. □

We continue with the orientation of the edges in G. In the following, we assume that given a node $\operatorname{enc}(\sigma) \in V$, we are able to compute in polynomial time the vertices of the corresponding simplex $\sigma \in \Sigma$. We show afterwards how to implement this step. With this assumption, the orientation can be defined similarly as in [23].

Let $\operatorname{enc}(\sigma), \operatorname{enc}(\sigma') \in V_d$ be two adjacent nodes. By definition, the encoded simplices $\sigma = \operatorname{conv}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_{d-1})$ and σ' share a facet $\check{\sigma} = \operatorname{conv}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-1})$ with $\lambda(\check{\sigma}) = [d-1]$. Let the indices be such that $\lambda(\boldsymbol{v}_i) = i$ for $i \in [d-1]$. Then, the edge between $\mathrm{enc}(\sigma)$ and $\mathrm{enc}(\sigma')$ is directed from $\mathrm{enc}(\sigma)$ to $\mathrm{enc}(\sigma')$ if and only if the function $\mathrm{dir}(\sigma, \sigma')$ is positive, where

$$\mathrm{dir}(\sigma, \sigma') = \operatorname{sgn}\det \begin{pmatrix} 1 & 1 & \dots & 1 \\ \boldsymbol{v}_0 & \boldsymbol{v}_1 & \dots & \boldsymbol{v}_{d-1} \end{pmatrix}.$$

Only for the sake of orientation, we define a set of $d-1$ vertices $\boldsymbol{w}_2, \dots, \boldsymbol{w}_d$ with colors $2, \dots, d$ to lift lower-dimensional simplices in order to avoid dealing with simplices of different dimensions. For $i = 2, \dots, d$, let $\boldsymbol{w}_i \in \mathbb{R}^d$ denote the parameter vector

$$(\boldsymbol{w})_j = \begin{cases} 2 & \text{if } j < i, \\ 1 - 2(i-1) & \text{if } j = i, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

where $j \in [d]$. Furthermore, we set $\lambda(\boldsymbol{w}_i) = i$. Since $(\boldsymbol{w}_i)_i < 0$ for $i = 2, \dots, d$, we have $\boldsymbol{w}_i \notin \Delta$ and for $k < i$, $\boldsymbol{w}_i \notin \mathrm{aff}(\Delta_{[k]})$. However, a quick calculation shows that $\boldsymbol{w}_i \in \mathrm{aff}(\Delta_{[i]})$ and that within $\mathrm{aff}(\Delta_{[i]})$, the hyperplane $\mathrm{aff}(\Delta_{[i-1]})$ separates $\boldsymbol{e}_i$ and $\boldsymbol{w}_i$. Now, let $\sigma = \mathrm{conv}(\boldsymbol{v}_0, \dots, \boldsymbol{v}_{k-1})$ denote a simplex that corresponds to some node in G, where $k \in [d-1]_0$. Then, we denote with $\sigma_{\boldsymbol{w}} = \mathrm{conv}(\boldsymbol{v}_0, \dots, \boldsymbol{v}_{k-1}, \boldsymbol{w}_{k+1}, \dots, \boldsymbol{w}_d)$ the $(d-1)$-simplex that we obtain by lifting σ with our additional vertices outside of Δ. Note that $\sigma_{\boldsymbol{w}}$ is non-degenerate by our choice of $\boldsymbol{w}_2, \dots, \boldsymbol{w}_d$. If σ is already a $(d-1)$-simplex, we set $\sigma_{\boldsymbol{w}} = \sigma$. Let now $\mathrm{enc}(\sigma)$ and $\mathrm{enc}(\sigma') \in V$ be two adjacent nodes. Then the two lifted simplices $\sigma_{\boldsymbol{w}}$ and $\sigma'_{\boldsymbol{w}}$ share a $[d-1]$-labeled facet. Now, we set $\mathrm{dir}(\sigma, \sigma') = \mathrm{dir}(\sigma_{\boldsymbol{w}}, \sigma'_{\boldsymbol{w}})$ and we direct the edge between $\mathrm{enc}(\sigma')$ and $\mathrm{enc}(\sigma)$ as discussed before. The following lemma guarantees that the orientation of the edge is the same if seen from either σ or σ' and that the only sinks and sources remain the nodes of degree 1 that are characterized by Lemma 4.8.

Lemma F.1. *The orientation of G is well-defined. Furthermore,* $\mathrm{enc}(\sigma) \in V$ *is a sink or a source if and only if* $\deg \mathrm{enc}(\sigma) = 1$ *in the underlying undirected graph.*

Proof. Let $\mathrm{enc}(\sigma), \mathrm{enc}(\sigma') \in V$ be two adjacent nodes. Assume first that $\mathrm{enc}(\sigma), \mathrm{enc}(\sigma') \in V_k$ for some $k \in [d]$. Let $\sigma = \mathrm{conv}(\boldsymbol{v}_0, \boldsymbol{v}_1, \dots, \boldsymbol{v}_{k-1})$ and $\sigma' = \mathrm{conv}(\boldsymbol{v}'_0, \boldsymbol{v}_1, \dots, \boldsymbol{v}_{k-1})$ denote the encoded simplices with $\lambda(\boldsymbol{v}_i) = i$ for $i \in [k-1]$. That is, σ and σ' share the facet $\check{\sigma} = \mathrm{conv}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_{k-1})$. Because both simplices are contained in Σ_k, the two vertices $\boldsymbol{v}_0$ and $\boldsymbol{v}'_0$ are separated within the $(k-1)$-dimensional affine space $\mathrm{aff}(\Delta_{[k]})$ by the $(k-2)$-dimensional affine space $\mathrm{aff}(\check{\sigma})$. Since $\boldsymbol{w}_{k+1}, \dots, \boldsymbol{w}_d \notin \mathrm{aff}(\Delta_{[k]})$, the two vertices $\boldsymbol{v}_0$ and $\boldsymbol{v}'_0$ are separated in $\mathbb{R}^d$ by the hyperplane $\mathrm{aff}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_{k-1}, \boldsymbol{w}_{k+1}, \dots, \boldsymbol{w}_d)$. Then, we have $\mathrm{dir}(\sigma, \sigma') = -\mathrm{dir}(\sigma', \sigma)$, since

$$\begin{aligned} \mathrm{dir}(\sigma, \sigma') &= \mathrm{dir}(\sigma_{\boldsymbol{w}}, \sigma'_{\boldsymbol{w}}) \\ &= \operatorname{sgn}\det \begin{pmatrix} 1 & 1 & \dots & 1 & 1 & \dots & 1 \\ \boldsymbol{v}_0 & \boldsymbol{v}_1 & \dots & \boldsymbol{v}_{k-1} & \boldsymbol{w}_{k+1} & \dots & \boldsymbol{w}_d \end{pmatrix} \\ &= -\operatorname{sgn}\det \begin{pmatrix} 1 & 1 & \dots & 1 & 1 & \dots & 1 \\ \boldsymbol{v}'_0 & \boldsymbol{v}_1 & \dots & \boldsymbol{v}_{k-1} & \boldsymbol{w}_{k+1} & \dots & \boldsymbol{w}_d \end{pmatrix} \\ &= -\mathrm{dir}(\sigma'_{\boldsymbol{w}}, \sigma_{\boldsymbol{w}}) \\ &= -\mathrm{dir}(\sigma', \sigma). \end{aligned}$$

Let now $\mathrm{enc}(\sigma) \in V_{k-1}$ and $\mathrm{enc}(\hat{\sigma}) \in V_k$ be two adjacent nodes for some $k \in [d]$. By definition of E, we then have $\lambda(\sigma) = [k-1]$ and σ is a facet of $\hat{\sigma}$. We write $\sigma = \mathrm{conv}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_{k-1})$ and

$\hat{\sigma} = \operatorname{conv}(\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_{k-1})$, where the indices are such that $\lambda(\boldsymbol{v}_i) = i$ for $i \in [k-1]$. Then,

$$\sigma_{\boldsymbol{w}} = \operatorname{conv}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{k-1}, \boldsymbol{w}_k, \ldots, \boldsymbol{w}_d) \text{ and } \hat{\sigma}_{\boldsymbol{w}} = \operatorname{conv}(\boldsymbol{v}_0, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_{k-1}, \boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_d).$$

Hence, $\sigma_{\boldsymbol{w}}$ and $\hat{\sigma}_{\boldsymbol{w}}$ share the facet $\check{\sigma}_{\boldsymbol{w}} = \operatorname{conv}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{k-1}, \boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_d)$. By construction, both vertices $\boldsymbol{v}_0$ and $\boldsymbol{w}_k$ are contained in $\operatorname{aff}(\Delta_{[k]})$. Within the $(k-1)$-dimensional affine space $\operatorname{aff}(\Delta_{[k]})$, the vertex $\boldsymbol{w}_k$ is separated from $\Delta_{[k]}$ by the $(k-2)$-dimensional affine space $\operatorname{aff}(\Delta_{[k-1]})$ and hence it is separated from $\boldsymbol{v}_0$ by $\operatorname{aff}(\Delta_{[k-1]})$. Since $\sigma \in \Sigma_{k-1}$, σ is a $(k-2)$ simplex that is contained in $\Delta_{[k-1]}$ and thus $\operatorname{aff}(\sigma) = \operatorname{aff}(\Delta_{[k-1]})$ separates $\boldsymbol{v}_0$ and $\boldsymbol{w}_k$ in $\operatorname{aff}(\Delta_{[k]})$. Now, because $\boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_k \notin \operatorname{aff}(\Delta_{[k]})$, $\boldsymbol{v}_0$ and $\boldsymbol{w}_k$ are separated in $\mathbb{R}^d$ by the hyperplane $\operatorname{aff}(\check{\sigma}_{\boldsymbol{w}})$. Again we have $\operatorname{dir}(\sigma, \hat{\sigma}) = -\operatorname{dir}(\hat{\sigma}, \sigma)$, since

$$\begin{aligned}
\operatorname{dir}(\sigma, \hat{\sigma}) &= \operatorname{dir}(\sigma_{\boldsymbol{w}}, \hat{\sigma}_{\boldsymbol{w}}) \\
&= \operatorname{sgn}\det\begin{pmatrix} 1 & 1 & \ldots & 1 & 1 & \ldots & 1 \\ \boldsymbol{w}_k & \boldsymbol{v}_1 & \ldots & \boldsymbol{v}_{k-1} & \boldsymbol{w}_{k+1} & \ldots & \boldsymbol{w}_d \end{pmatrix} \\
&= -\operatorname{sgn}\det\begin{pmatrix} 1 & 1 & \ldots & 1 & 1 & \ldots & 1 \\ \boldsymbol{v}_0 & \boldsymbol{v}_1 & \ldots & \boldsymbol{v}_{k-1} & \boldsymbol{w}_{k+1} & \ldots & \boldsymbol{w}_d \end{pmatrix} \\
&= -\operatorname{dir}(\hat{\sigma}_{\boldsymbol{w}}, \sigma_{\boldsymbol{w}}) \\
&= -\operatorname{dir}(\hat{\sigma}, \sigma).
\end{aligned}$$

It remains to show the second part of the statement. Let $\operatorname{enc}(\sigma) \in V$ be a node with two adjacent nodes $\operatorname{enc}(\sigma'), \operatorname{enc}(\sigma'')$. We want to show that the two incident edges are oriented differently. In any case, the lifted simplices $\sigma_{\boldsymbol{w}}$ and $\sigma'_{\boldsymbol{w}}$ share a $[d-1]$-labeled facet $\check{\sigma}'_{\boldsymbol{w}}$ and similarly, $\sigma_{\boldsymbol{w}}$ and $\sigma''_{\boldsymbol{w}}$ share a $[d-1]$-labeled facet $\check{\sigma}''_{\boldsymbol{w}}$. The facets $\check{\sigma}'_{\boldsymbol{w}}$ and $\check{\sigma}''_{\boldsymbol{w}}$ of $\sigma_{\boldsymbol{w}}$ differ in exactly one vertex with the same label. Thus, the determinants in $\operatorname{dir}(\sigma, \sigma')$ and $\operatorname{dir}(\sigma, \sigma'')$ differ by exactly one column-swap. The properties of the determinant now ensure that $\operatorname{dir}(\sigma, \sigma') = -\operatorname{dir}(\sigma, \sigma'')$, as desired. □

Our next lemma shows that for purposes of orientation, we can replace the barycenters by arbitrary interior points in the corresponding parameter faces.

Lemma F.2. *Let $q_0, \ldots, q_{k-1} \subset \mathbb{R}^d$ be k polytopes such that $q_0 \subset \cdots \subset q_{k-1}$ and $\dim q_i = i$ for $i \in [k-1]_0$. Furthermore let $\boldsymbol{v}_i$ denote the barycenter of q_i for $i \in [k-1]_0$ and let $\boldsymbol{v}'_0, \ldots, \boldsymbol{v}'_{k-1}$ be $k-1$ vectors such that $\boldsymbol{v}'_i \in q_i$ and $\operatorname{aff}(\boldsymbol{v}'_0, \ldots, \boldsymbol{v}'_i) = \operatorname{aff}(q_i)$ for all $i \in [k-1]_0$. Then,*

$$\operatorname{sgn}\det\begin{pmatrix} 1 & \ldots & 1 & 1 & \ldots & 1 \\ \boldsymbol{v}_0 & \ldots & \boldsymbol{v}_{k-1} & \boldsymbol{x}_{k+1} & \ldots & \boldsymbol{x}_d \end{pmatrix} = \operatorname{sgn}\det\begin{pmatrix} 1 & \ldots & 1 & 1 & \ldots & 1 \\ \boldsymbol{v}'_0 & \ldots & \boldsymbol{v}'_{k-1} & \boldsymbol{x}_{k+1} & \ldots & \boldsymbol{x}_d \end{pmatrix},$$

where $\boldsymbol{x}_i \in \mathbb{R}^d \setminus \operatorname{aff} q_{k-1}$, $i \in [d] \setminus [k]$, is an arbitrary point.

Proof. The prove involves only basic linear algebra, however it is included for completeness. We show by induction on i that $\operatorname{aff}(q_i) = \operatorname{aff}(\boldsymbol{v}'_0, \ldots, \boldsymbol{v}'_i)$ and that for all $j \in [i]_0$, $\boldsymbol{v}'_j = \sum_{l=0}^{j} \alpha_{j,l} \boldsymbol{v}_l$ is an affine combination of $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_j$ with $\alpha_{j,j} > 0$.

For $i = 0$ the induction hypothesis trivially holds since $\dim q_0 = 0$ and hence $q_0 = \boldsymbol{v}_0 = \boldsymbol{v}'_0$. Assume now that $i > 0$ and that the induction hypothesis holds for all $i' < i$. Since q_{i-1} is a facet of q_i, within the i-dimensional affine space $\operatorname{aff}(q_i)$, q_i lies on one side of the $(i-1)$-dimensional affine space $\operatorname{aff}(q_{i-1})$ and thus it lies on one side of $\operatorname{aff}(\boldsymbol{v}'_0, \ldots, \boldsymbol{v}'_{i-1})$. Since both $\boldsymbol{v}_i$ and $\boldsymbol{v}'_i$ lie on the same side of $\operatorname{aff}(\boldsymbol{v}'_0, \ldots, \boldsymbol{v}'_{i-1})$ in $\operatorname{aff}(q_i)$, we can write $\boldsymbol{v}'_i$ as $\sum_{l=0}^{i-1} \beta_l \boldsymbol{v}'_l + \alpha_i \boldsymbol{v}_i$ with $\alpha_i > 0$. By our induction hypothesis, $\boldsymbol{v}'_0, \ldots, \boldsymbol{v}'_{i-1} \in \operatorname{aff}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_{i-1})$ and hence the hypothesis holds for i. The claim now follows directly from the properties of the determinant:

$$\begin{aligned}
&\operatorname{sgn}\det\begin{pmatrix} 1 & \dots & 1 & \dots & 1 & 1 & \dots & 1 \\ \boldsymbol{v}'_0 & \dots & \boldsymbol{v}'_i & \dots & \boldsymbol{v}'_{k-1} & \boldsymbol{x}_{k+1} & \dots & \boldsymbol{x}_d \end{pmatrix} \\
=&\operatorname{sgn}\det\begin{pmatrix} 1 & \dots & 1 & \dots & 1 & 1 & \dots & 1 \\ \boldsymbol{v}_0 & \dots & \sum_{l=0}^{i}\alpha_{i,l}\boldsymbol{v}_l & \dots & \sum_{l=0}^{k-1}\alpha_{k-1,l}\boldsymbol{v}_l & \boldsymbol{x}_{k+1} & \dots & \boldsymbol{x}_d \end{pmatrix} \\
=&\operatorname{sgn}\det\begin{pmatrix} 1 & \dots & 1 & \dots & 1 & 1 & \dots & 1 \\ \boldsymbol{v}_0 & \dots & \alpha_{i,i}\boldsymbol{v}_l & \dots & \alpha_{k-1,k-1}\boldsymbol{v}_{k-1} & \boldsymbol{x}_{k+1} & \dots & \boldsymbol{x}_d \end{pmatrix} \\
=&\operatorname{sgn}\det\begin{pmatrix} 1 & \dots & 1 & \dots & 1 & 1 & \dots & 1 \\ \boldsymbol{v}_0 & \dots & \boldsymbol{v}_i & \dots & \boldsymbol{v}_{k-1} & \boldsymbol{x}_{k+1} & \dots & \boldsymbol{x}_d \end{pmatrix},
\end{aligned}$$

where the last equality holds since $\alpha_{i,i} > 0$ for $i \in [k-1]$. □

As the next lemma shows, computing parameter vectors in the relative interior of faces in $\mathcal{Q}_\Delta$ is computationally feasible.

Lemma F.3. *Let* $\operatorname{enc}(\sigma) = (\operatorname{enc}(q_0), \dots, \operatorname{enc}(q_{k-1})) \in V$ *be a node of* G*, where* $k \in [d]$*. Then, we can compute in polynomial time* $k-1$ *parameter vectors* $\boldsymbol{v}_0, \dots, \boldsymbol{v}_{k-1}$ *such that* $\boldsymbol{v}_i \in q_i$ *and* $\operatorname{aff}(\boldsymbol{v}_0, \dots, \boldsymbol{v}_i) = \operatorname{aff}(q_i)$ *for* $i \in [k-1]_0$*.*

Proof. By definition of the encoding, q_0 is a vertex and hence we can choose $\boldsymbol{v}_0 = q_0$. The algorithm iteratively computes now incident edges $e_i = \operatorname{conv}(\boldsymbol{v}_0, \boldsymbol{v}_i)$ to $\boldsymbol{v}_0$ for $i \in [k-1]$ such that e_i is an edge of q_i and no edge of q_{i-1}. The resulting vectors have the desired properties: $\boldsymbol{v}_i \in q_i$ and $\operatorname{aff}(\boldsymbol{v}_0, \dots, \boldsymbol{v}_i) = \operatorname{aff}(q_i)$ for $i \in [k-1]_0$.

We construct these edges as follows. Write $\operatorname{enc}(q)_i = \left(\operatorname{supp}(f_i), I_0^{(i)}, I_1^{(i)}\right)$ and let g_i be the face $g\left(I_0^{(i)}, I_1^{(i)}\right)$ of $\mathcal{M}$ that is encoded by the index sets $I_0^{(i)}$ and $I_1^{(i)}$. Since $\operatorname{enc}(\sigma)$ is a valid k-tuple, the columns B from $A_{\operatorname{supp}(f_{k-1})}$ are a feasible basis and moreover, since $\operatorname{supp}(f_{k-1}) \subseteq \operatorname{supp}(f_i)$ for $i \in [k-1]_0$, the set B is a feasible basis for all faces f_i, $i \in [k-1]_0$. Similar to the proof of Lemma 4.6, we can express each polytope $\mathcal{M}(q_i)$ as the solution to the linear system L^{Φ}_{B,f_i} extended by the constraints $\boldsymbol{\mu} \in g_i$, where $i \in [k-1]_0$. Let L_i denote the resulting linear system. Again by the properties of a valid k-tuple, either $\operatorname{supp}(f_{i-1}) = \operatorname{supp}(f_i) \cup \{a_{i-1}\}$, where $a_i \in [d^2] \setminus \operatorname{supp}(f_i)$. Or there is an index $j_{i-1} \in [d] \setminus \left(I_0^{(i)} \cup I_1^{(i)}\right)$ such that $I_0^{(i-1)} = I_0^{(i)} \cup \{j_{i-1}\}$ and $I_1^{(i-1)} = I_1^{(i)}$, or $I_0^{(i-1)} = I_0^{(i)}$ and $I_1^{(i-1)} = I_1^{(i)} \cup \{j_{i-1}\}$. This means, that the linear system L_{i-1} equals the linear system L_i where one inequality becomes tight. In the following we call this inequality e_i. Note that L_0 is then the linear system L_{k-1} in which all inequalities $e_1, \dots, e_{k-1}$ are tight.

Assume now that we already have computed the vectors $\boldsymbol{v}_0, \dots, \boldsymbol{v}_{i-1}$ such that $\boldsymbol{v}_j \in q_j$ and $\operatorname{aff}(\boldsymbol{v}_0, \dots, \boldsymbol{v}_j) = \operatorname{aff}(q_j)$ for $j \in [i-1]_0$ and we want to compute $\boldsymbol{v}_i$, where $i \in [k-1]$. We consider the linear system L'_i that we obtain by relaxing the tight inequality e_i in L_0. Since the solution space of L_0 is the vertex $\boldsymbol{v}_0$, the solution space to L'_i is an edge $\operatorname{conv}(\boldsymbol{v}_0, \boldsymbol{v}_i)$. We can compute the other endpoint $\boldsymbol{v}_i$ of this edge in polynomial time by computing the line that is defined by the equalities in L'_i and intersect this iteratively with the halfspaces that are defined by the inequalities in L'_i while keeping track of the endpoints. Now, we have $\boldsymbol{v}_i \in q_i$ since the solution space of the linear system L'_i is a subset of the solution space of the linear system L_i. Moreover, since in L_{i-1} the inequality e_i is tight, $\boldsymbol{v}_i \in q_i \setminus q_{i-1}$ and thus $\operatorname{aff}(\boldsymbol{v}_0, \dots, \boldsymbol{v}_i) = \operatorname{aff}(q_i)$. □

The following lemma is now an immediate consequence of Lemmas F.2 and F.3.

Lemma F.4. *Let* $\operatorname{enc}(\sigma), \operatorname{enc}(\sigma) \in V$ *be two adjacent nodes. Then, we can compute* $\operatorname{dir}(\sigma, \sigma')$ *in polynomial time.* □

G A Polynomial-Time Case

In the following, we use the same notation as in Section 3 (see Table 1 on Page 14 for an overview). Let $C_1, C_2 \subset \mathbb{Q}^d$ be two color classes, each of size d, let $\boldsymbol{b} \in \mathbb{Q}^d$, $\boldsymbol{b} \neq \boldsymbol{0}$, be a point that is ray-embraced by C_1 and by C_2, and let $k \in [d-1]$ be a number. Although not needed in the algorithm, to comply with the formulations of our results in Section B and Section 3, we introduce $d-2$ "dummy" color classes $C_3, \dots, C_d$ that trivially ray-embrace $\boldsymbol{b}$ by setting $C_3 = \dots = C_d = \{\boldsymbol{b}\}$. Let $(C'_1, \dots, C'_d, \boldsymbol{b}')$ be the instance of ColorfulCarathéodory in general position that we obtain by applying Lemma B.6 to $(C_1, \dots, C_d, \boldsymbol{b})$. Then, let $\mathcal{P}^{\mathrm{CC}} \subset \mathbb{Q}^{d^2}$ denote the polyhedron that is defined by the linear system L^{CC} (see (2) on Page 6) for the instance $(C'_1, \dots, C'_d, \boldsymbol{b}')$. Furthermore, let $\Delta_1 = \Delta \cap \operatorname{conv}(\boldsymbol{e}_1, \boldsymbol{e}_2)$ denote the edge of the standard simplex Δ^{d-1} that connects $\boldsymbol{e}_1$ with $\boldsymbol{e}_2$ and set $\mathcal{Q}_{\Delta_1} = \{q \in \mathcal{Q}_\Delta \mid q \subseteq \Delta_1\}$. Note that by Lemma 4.3, the set Q_{Δ_1} is a 1-dimensional polytopal complex that decomposes Δ_1. We begin with the following basic lemma on $\mathcal{Q}_{\Delta_1}$.

Lemma G.1. *Let $e, e' \in \mathcal{Q}_{\Delta_1}$, $e \neq e'$, be two adjacent edges with $e = \Phi_\Delta(f) \cap g$ and $e' = \Phi_\Delta(f') \cap g'$, where $f, f' \in \mathcal{F}$ and $g, g' \in \mathcal{S}$. Then, f and f' are vertices of $\mathcal{P}^{\mathrm{CC}}$ with* $\operatorname{supp}(f), \operatorname{supp}(f') \subseteq \operatorname{ind}(C'_1 \cup C'_2)$ *and* $\operatorname{supp}(f), \operatorname{supp}(f')$ *differ in at most one column index.*

Proof. By Lemma 4.2, the faces f, f' are vertices of $\mathcal{P}^{\mathrm{CC}}$. Furthermore, since $\mathcal{M}(e), \mathcal{M}(e') \subset \operatorname{span}(\boldsymbol{e}_1, \boldsymbol{e}_2)$, Lemma 4.1 implies that $\operatorname{supp}(f), \operatorname{supp}(f') \subseteq \operatorname{ind}(C'_1 \cup C'_2)$. Now, since e and e' are adjacent, they share a vertex $\boldsymbol{v} = \Phi_\Delta(f_{\boldsymbol{v}}) \cap g_{\boldsymbol{v}} \in \mathcal{Q}_{\Delta_1}$, where $f_{\boldsymbol{v}} \in \mathcal{F}$ and $g_{\boldsymbol{v}} \in \mathcal{S}$. Then, by Lemma 4.2, either f is a facet of $f_{\boldsymbol{v}}$ and $g = g_{\boldsymbol{v}}$, or $f = f_{\boldsymbol{v}}$ and $g_{\boldsymbol{v}}$ is a facet of g. Similarly, either f' is a facet of $f_{\boldsymbol{v}}$ and $g' = g_{\boldsymbol{v}}$, or $f' = f_{\boldsymbol{v}}$ and $g_{\boldsymbol{v}}$ is a facet of g'. Then, Observation D.7 implies the statement. □

Using Lemma G.1, we now present a polynomial-time checkable criterion whether an interval $[\boldsymbol{\mu}_1, \boldsymbol{\mu}_2] \subset \Delta_1$ intersects an edge $e^\star = \Phi_\Delta(f^\star) \cap g^\star \in \mathcal{Q}_{\Delta_1}$, where $f \in \mathcal{F}$ and $g \in \mathcal{S}$, such that $\operatorname{supp}(f^\star)$ defines a $(k, d-k)$-colorful choice that ray-embraces $\boldsymbol{b}'$.

Corollary G.2. *Let $k \in [d-1]$, be a number and let $e, e' \in \mathcal{Q}_{\Delta_1}$ be two edges with $e = \Phi_\Delta(f) \cap g$ and $e' = \Phi_\Delta(f') \cap g'$, where $f, f' \in \mathcal{F}$ and $g, g' \in \mathcal{S}$. If $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f)| < k$ and $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f')| > k$, then there exists an edge $e^\star = \Phi_\Delta(f^\star) \cap g^\star \subset \operatorname{conv}(e, e')$, $e^\star \in \mathcal{Q}_{\Delta_1}$, such that $\operatorname{supp}(f^\star)$ defines a $(k, d-k)$-colorful choice of C_1 and C_2 that ray-embraces $\boldsymbol{b}'$, where $f^\star \in \mathcal{F}$ and $g^\star \in \mathcal{S}$.*

Proof. By Lemma G.1, the supports of the faces in $\mathcal{F}$ that corresponds to two adjacent edges in $\mathcal{Q}_{\Delta_1}$ differ in at most one column. Since $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f)| < k$, $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f')| > k$, and since Q_{Δ_1} is a polytopal complex, there must be an edge $e^\star = \Phi_\Delta(f^\star) \cap g^\star \in \mathcal{Q}_{\Delta_1}$ between e and e' such that $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f^\star)| = k$. By Lemma 4.2, $f^\star$ is a vertex and hence $|\operatorname{supp}(f^\star)| = d$. In particular, then $|\operatorname{ind}(C_2) \cap \operatorname{supp}(f^\star)| = d - k$. □

The algorithm to find this $(k, d-k)$-colorful choice is now a straightforward application of binary search. Initially we set $\boldsymbol{\mu}_1 = \boldsymbol{e}_1$ and $\boldsymbol{\mu}_2 = \boldsymbol{e}_2$ and we maintain the invariant that the interval $[\boldsymbol{\mu}_1, \boldsymbol{\mu}_2]$ contains an edge $e^\star = \Phi_\Delta(f^\star) \cap g^\star \in \mathcal{Q}_{\Delta_1}$ such that $\operatorname{supp}(f^\star)$ defines a $(k, d-k)$-colorful choice that ray-embraces $\boldsymbol{b}'$. The single optimal feasible basis for $\boldsymbol{e}_1$ is C_1 and similarly, the single optimal feasible basis for $\boldsymbol{e}_2$ is C_2. Then, Corollary G.2 implies the invariant for the initial interval. We repeatedly proceed as follows: set $\boldsymbol{\mu}' = \frac{1}{2}(\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2)$ and solve the linear program $L^{\mathrm{CC}}_{\mathcal{M}(\boldsymbol{\mu}')}$. Let $\operatorname{supp}(f')$ be the support of the maximum face $f' \in \mathcal{F}$ that is optimal for $L^{\mathrm{CC}}_{\mathcal{M}(\boldsymbol{\mu}')}$. First assume that $|\operatorname{supp}(f')| = d$, i.e., assume that f' is a vertex of $\mathcal{P}^{\mathrm{CC}}$. If $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f')| = k$, we have found the desired solution. If $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f')| < k$, we set $\boldsymbol{\mu}_2 = \boldsymbol{\mu}'$ and otherwise, if $|\operatorname{ind}(C_1) \cap \operatorname{supp}(f')| > k$, we set $\boldsymbol{\mu}_1 = \boldsymbol{\mu}'$. By Corollary G.2, the invariant is maintained. Now, assume that $|\operatorname{supp}(f')| = d+1$, i.e., assume that f' is an edge of $\mathcal{P}^{\mathrm{CC}}$. Then, by Lemma 4.2, $\boldsymbol{\mu}' = \Phi_\Delta(f') \cap g$ is a vertex of $\mathcal{Q}_{\Delta_1}$ and since $\boldsymbol{\mu}' \in \operatorname{relint} \Delta_1$, it is incident to two edges $e_1, e_2 \in \mathcal{Q}_{\Delta_1}$ with $e_1 = \Phi_\Delta(f_1) \cap g$ and $e_2 = \Phi_\Delta(f_2) \cap g$, where f_1 and f_2 are the two incident vertices to the edge f'. We compute both supports $\operatorname{supp}(f_1)$ and $\operatorname{supp}(f_2)$ by checking every d-subset of $\operatorname{supp}(f')$ whether it constitutes a basis. Then, we check whether one of the two supports is a $(k, d-k)$-colorful choice. If not, then by Lemma G.1, either both supports contain less than k columns from C_1 or both contain more than k columns from C_1. In the first case, we set $\boldsymbol{\mu}_2 = \boldsymbol{\mu}'$ and in the second case, we set $\boldsymbol{\mu}_1 = \boldsymbol{\mu}'$. Again, Corollary G.2 guarantees that the invariant is maintained.

Clearly, each update of the interval $[\boldsymbol{\mu}_1, \boldsymbol{\mu}_2]$ needs weakly polynomial time since $O(d)$ linear programs are solved. Furthermore, the number of the steps needed before a solution is found is logarithmic in the length of the shortest edge. The following lemma shows that the minimum length of an edge in $\mathcal{Q}_{\Delta_1}$ is at least exponentially small in the length of the ColorfulCarathéodory instance.

Lemma G.3. *Let L be the length of the binary encoding of the* ColorfulCarathéodory *instance $(C'_1, \ldots, C'_d, \boldsymbol{b}')$ and let $e = [\boldsymbol{\mu}_1, \boldsymbol{\mu}_2] \in \mathcal{Q}_{\Delta_1}$ be an edge. Then, $-\log \|\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1\| = \Omega(\operatorname{poly} L)$.*

Proof. We write e as $\Phi_\Delta(f) \cap g$ and the two incident vertices as $\boldsymbol{\mu}_1 = \Phi_\Delta(f_1) \cap g_1$ and $\boldsymbol{\mu}_2 = \Phi_\Delta(f_2) \cap g_2$, where $\{f, f_1, f_2\} \subseteq \mathcal{F}$ and $\{g, g_1, g_2\} \subseteq \mathcal{S}$. We denote with $\hat{\boldsymbol{\mu}}_1 = \mathcal{M}(\boldsymbol{\mu}_1)$ and with $\hat{\boldsymbol{\mu}}_1 = \mathcal{M}(\boldsymbol{\mu}_1)$ the vertices in $\mathcal{Q}$ whose central projections onto Δ resulted in $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$, respectively. Since e is an edge, $\hat{\boldsymbol{\mu}}_1 \neq \hat{\boldsymbol{\mu}}_2$ and hence there is a $j \in [d]$ with $(\hat{\boldsymbol{\mu}}_1)_j \neq (\hat{\boldsymbol{\mu}}_2)_j$. By Lemma 4.2, f is a vertex of $\mathcal{P}^{\mathrm{CC}}$ and $\operatorname{supp}(f) \subseteq \operatorname{supp}(f_i)$ for $i = 1, 2$. Let B denote the columns in $A_{\operatorname{supp}(f)}$. Then, we can express $\hat{\boldsymbol{\mu}}_i$, $i = 1, 2$, as the unique solution to the linear system L^{Φ}_{B, f_i} extended by the constraints $\boldsymbol{\mu} \in \mathcal{M}(g_i)$. Now, Lemma D.1 guarantees that the logarithm of $(\hat{\boldsymbol{\mu}}_i)_j$, $i \in [2]$, is a polynomial in the size of the linear system and hence in L. Since $(\boldsymbol{\mu}_1)_j \neq (\boldsymbol{\mu}_2)_j$, we have $= -\log \|\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1\| = \Omega(\operatorname{poly} L)$, as claimed. □

The described binary-search algorithm needs therefore only polynomial time in L to compute a $(k, d-k)$-colorful choice C' for C'_1 and C'_2. Since L is polynomial in the length of the of the original instance $(C_1, \ldots, C_d, \boldsymbol{b})$, the running time is weakly polynomial in the length of the original instance. Furthermore, we can obtain a $(k, d-k)$-colorful choice C for C_1 and C_2 by replacing the perturbed points in C' with the original points in $C_1 \cup C_2$. Lemma B.5 then guarantees that C ray-embraces $\boldsymbol{b}$.