The Colorful Carathéodory Problem and its Descendants

Dissertation zur Erlangung des Doktorgrades

vorgelegt am

Fachbereich Mathematik und Informatik
der Freien Universität Berlin

2016

von

Yannik Stein

Betreuer: Prof. Dr. Wolfgang Mulzer
Zweitgutachter: Prof. Dr. Donald R. Sheehy
Datum der Disputation: 21. Oktober 2016

Abstract

The colorful Carathéodory theorem is an existence theorem that implies several statements on convex intersection patterns such as Tverberg's theorem, the centerpoint theorem, the first selection lemma, and the colorful Kirchberger theorem. Interestingly, these proofs can be interpreted as polynomial-time reductions to COLORFULCARATHÉODORY, the computational search problem that corresponds to the colorful Carathéodory theorem. We exploit this existing web of reductions by developing approximation algorithms and complexity bounds on COLORFULCARATHÉODORY that also apply to its polynomial-time descendants.

Let $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ be finite point sets such that $\mathbf{0} \in \operatorname{conv}(C_i)$ for $i \in [d+1]$. Then, the colorful Carathéodory theorem asserts that we can choose one point from each set C_i such that the chosen points C contain the origin in their convex hull. COLORFULCARATHÉODORY is then the computational problem of finding C. Since a solution always exists and since it can be verified in polynomial time, COLORFULCARATHÉODORY is contained in *total function NP* (TFNP), the class of NP search problems that always admit a solution. We show that COLORFULCARATHÉODORY belongs to the intersection of two important subclasses of TFNP: the complexity classes *polynomial-time parity argument on directed graphs* (PPAD) and *polynomial-time local search* (PLS). The formulation of COLORFULCARATHÉODORY as a PPAD-problem is based on a new constructive proof of the colorful Carathéodory theorem that uses Sperner's lemma. Moreover, we show that already a slight change in the definition of COLORFULCARATHÉODORY results in a PLS-complete problem.

In the second part, we present several constructive results. First, we consider an approximation version of COLORFULCARATHÉODORY in which we are allowed to take more than one point from each set C_i. This notion of approximation has not been studied before and it is compatible with the polynomial-time reductions to COLORFULCARATHÉODORY. For any fixed $\varepsilon > 0$, we can compute a set C with $\mathbf{0} \in \operatorname{conv}(C)$ and at most $\lceil \varepsilon d \rceil$ points from each C_i in $d^{O(\varepsilon^{-1} \log \varepsilon^{-1})}$ time by repeatedly combining recursively computed approximations for lower-dimensional problem instances. Additionally, we consider a further notion of approximation in which we are given only $k < d+1$ sets C_i with $\mathbf{0} \in \operatorname{conv}(C_i)$, and we want to find a set C with at most $\lceil (d+1)/k \rceil$ points from each set C_i. The existence of C is a direct implication of the colorful Carathéodory theorem. Using linear programming techniques, we can solve the case $k = 2$ in weakly polynomial time. Moreover, we show that COLORFULCARATHÉODORY can be solved exactly in quasi-polynomial time when given $\operatorname{poly}(d)$ sets C_i that contain the origin in their convex hulls instead of only $d+1$. Finally, we consider the problem of computing the simplicial depth. The *simplicial depth* $\sigma_P(\boldsymbol{q})$ of a point $\boldsymbol{q} \in \mathbb{R}^d$ w.r.t. a set P is the number of distinct d-simplices with vertices in P that contain $\boldsymbol{q}$. If the dimension is constant, we show that $\sigma_P(\boldsymbol{q})$ can be $(1+\varepsilon)$-approximated w.h.p. in time $\widetilde{O}\left(n^{d/2+1}\right)$, where $\varepsilon > 0$ is an arbitrary constant. Furthermore, we show that the problem becomes #P-complete and W[1]-hard if the dimension is part of the input.

Acknowledgements

Most of all, I would like to thank my advisor Wolfgang Mulzer whose overwhelming support, patience, and guidance exceeded by far any expectations and made this thesis possible in the first place.

Further, I would like to thank Frédéric Meunier and Pauline Sarrabezolles for many interesting discussions on the colorful Carathéodory theorem, for the fruitful collaboration, and for teaching me much about linear programming. I am also very grateful for the nice and productive month that I spend at MADALGO in Aarhus and in particular I would like to thank Peyman Afshani for hosting me and for introducing me to simplicial depth. Moreover, I would like to thank Donald R. Sheehy for agreeing to review this thesis.

I wish to thank the whole group "Theoretische Informatik" at Freie Universität Berlin for providing a friendly working environment and in particular, I would like to thank: Paul Seiferth with whom it was always fun to work, my former officemate Rafel Jaume Deyà for great advice, my current officemate Bahar Banyassadi for a happy atmosphere in the office, and Frank Hoffmann for valuable feedback.

I also would like to thank Günter M. Ziegler for being my BMS mentor and valuable career advice. I gratefully acknowledge the scholarship from the RTG "Methods for Discrete Structures" and their more than generous travel support, as well as the support from GIF grant 1161.

Finally, I would like to thank my friends for their great support during my PhD and for always being a welcome source of distraction.

Contents

Abstract **v**

List of Figures **xi**

List of Algorithms **xiii**

1. Introduction **1**
1.1. Notation and Computational Model 2
1.2. Descendants of the Colorful Carathéodory Theorem 2
1.3. Previous Work 15

I. The Complexity of ColorfulCarathéodory **17**

2. Introduction to Part I **19**
2.1. Overview 21
2.2. Preliminaries 21
2.2.1. FNP Reductions 21
2.2.2. The Complexity Class PPAD 22
2.2.3. The Complexity Class PLS 23
2.2.4. Linear Programming 25

3. Equivalent Instances of ColorfulCarathéodory in General Position **29**
3.1. Polynomials with Bounded Integer Coefficients 29
3.2. Construction 32

4. ColorfulCarathéodory is in PPAD **35**
4.1. Outline 38
4.2. The Polytopal Complex 39
4.3. The Barycentric Subdivision 45
4.4. The PPAD graph 51

5. ColorfulCarathéodory is in PLS **57**
5.1. A PLS Formulation of ColorfulCarathéodory 57
5.2. A PLS-Complete Generalization of ColorfulCarathéodory 60

II. Algorithms for ColorfulCarathéodory **67**

6. Introduction to Part II **69**
6.1. Overview . . . 70

7. k-Colorful Choices **73**
7.1. Embracing Equivalent Points . . . 73
7.2. Computing k-Colorful Choices . . . 77
7.3. Applications . . . 83

8. Few Color Classes **87**

9. Exact Algorithms for ColorfulCarathéodory **91**
9.1. A Simple Special Case . . . 91
9.2. Many Colors . . . 92

10. Simplicial Depth **95**
10.1. Preliminaries . . . 95
10.1.1. The Gale Transform . . . 95
10.1.2. The Complexity Class #P . . . 98
10.1.3. The Complexity Classes FPT and W[1] . . . 98
10.2. Approximation in High Dimensions . . . 99
10.3. Improving the Brute-Force Approach . . . 101
10.4. Complexity . . . 103

11. Conclusions **107**

Bibliography **109**

List of Figures

1.1. Example of the colorful Carathéodory theorem in two dimensions. 4
1.2. Proof of the colorful Carathéodory theorem. 5
1.3. Approximating convex sets by halfspaces. 7
1.4. Example of Sarkaria's Lemma. 11
1.5. Implications of the colorful Carathéodory theorem. 16
1.6. Descendants of COLORFULCARATHÉODORY. 16

2.1. Example of a PPAD-graph problem. 22
2.2. Example of a PLS-graph for a minimization problem. 24

4.1. Example of Sperner's lemma in two dimensions. 36
4.2. Orienting the PPAD graph. 53

5.1. Construction of the L-NCP instance that corresponds to a given MAX-2SAT/FLIP instance. 63

7.1. Minimal **0**-embracing sets. 73
7.2. Embracing equivalent point sets and representative points. 75

10.1. Overview of the notation for the Gale transform. 96
10.2. Dimension-reduction argument in the improved brute-force algorithm to compute the simplicial depth. 102
10.3. Computing the simplicial depth is #P-complete and W[1]-hard. 105

List of Algorithms

7.1. Computing **0**-embracing $(\lceil d/2 \rceil + 1)$-colorful choices. 77
7.2. Computing **0**-embracing $\lceil \varepsilon d \rceil$-colorful choices. 79

9.1. Solving COLORFULCARATHÉODORY with many color classes exactly. 93

10.1. Improved brute-force algorithm to compute the simplicial depth. 103

1 Introduction

Centerpoints are a generalization of the median to higher dimensions with many applications in statistics [17]. Let $P \subset \mathbb{R}^d$ be a point set. We say a point $\boldsymbol{c} \in \mathbb{R}^d$ has *Tukey depth* [75] $\tau \in \mathbb{N}$ with respect to $P \subset \mathbb{R}^d$ if all closed halfspaces that contain $\boldsymbol{c}$ also contain at least τ points from P. The centerpoint theorem [65] guarantees that there always exists a point with Tukey depth $\lceil |P|/(d+1) \rceil$ and we call such a point a *centerpoint* for P. Although the computational problem of finding centerpoints (CENTERPOINT) is well-studied, its complexity remains elusive.

In general it is coNP-complete [74, Theorem 8.8] to verify whether a given point is a centerpoint. However, certain centerpoints have a polynomial-time checkable certificate. A *Tverberg m-partition* of P is a partition into $m \in \mathbb{N}$ sets whose convex hulls have a nonempty intersection. Tverberg's theorem [76] guarantees the existence of a Tverberg $(|P|/(d+1))$-partition and we call a point in the intersection of such a partition a *Tverberg point*. A Tverberg point $\boldsymbol{c}$ is a centerpoint since every halfspace that contains $\boldsymbol{c}$ has to contain at least one point from each element of the corresponding Tverberg $(|P|/(d+1))$-partition. The complexity of finding Tverberg $(|P|/(d+1))$-partitions (TVERBERG) is not settled.

Related to the Tukey depth, the *simplicial depth* is a further notion of data depth. The simplicial depth of a point $\boldsymbol{q} \in \mathbb{R}^d$ with respect to a set $P \subset \mathbb{R}^d$ is the number of distinct d-simplices with vertices in P that contain $\boldsymbol{q}$. The first selection lemma guarantees that there always exists a point with simplicial depth $\Omega\left(f(d)\, n^{d+1}\right)$, where $f : \mathbb{N} \mapsto \mathbb{R}_+$ is a function. If d is constant, this is asymptotically the best possible. The complexity of finding a point with simplicial depth $\Omega\left(f'(d)\, n^{d+1}\right)$ (SIMPLICIALCENTER), where $f' : \mathbb{N} \mapsto \mathbb{R}_+$ is an arbitrary but fixed function, is unknown.

All three computational problems, finding centerpoints, finding Tverberg partitions, and finding points with large simplicial depth, have in common that no efficient algorithms are known and only related computational problems have been properly categorized into the family of complexity classes. Interestingly, by interpreting Sarkaria's proof of Tverberg's theorem [69] algorithmically, it can be easily seen that all three problems are polynomial-time reducible to COLORFULCARATHÉODORY: given $d+1$ sets $C_1, \ldots, C_{d+1} \subset \mathbb{Q}^d$ that all contain the origin in their convex hulls, we want to compute a set C that contains one point from each C_i and that also contains the origin in its convex hull. The colorful Carathéodory theorem guarantees the existence of such a set C. The focus of this thesis is to derive new upper bounds on the complexity of COLORFULCARATHÉODORY and hence for CENTERPOINT, TVERBERG and SIMPLICIALCENTER. On the constructive side, we want to develop approximation algorithms for COLORFULCARATHÉODORY that are compatible with the reductions from CENTERPOINT, TVERBERG, and SIMPLICIALCENTER in the sense that a precise enough approximation of COLORFULCARATHÉODORY results in approximation algorithms for the others.

1.1. Notation and Computational Model

We denote with

- $\mathbb{R}_+ = \{x \in \mathbb{R} \mid x \geq 0\}$ the set of nonnegative reals;
- $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ the natural numbers with 0;
- $[k]$ the set $\{1, \ldots, k\}$ and with $[k]_0 = [k] \cup \{0\}$.

Throughout this thesis, symbols of vectors or points are set in boldface. The origin is denoted by $\mathbf{0}$, the canonical basis of $\mathbb{R}^d$ is denoted by $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d$, and the all-ones vector $\sum_{i=1}^d e_i$ is denoted by $\mathbf{1}$. For a set of points $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\} \subset \mathbb{R}^d$, we denote by

- $\operatorname{span}(P) = \left\{\sum_{i=1}^n \phi_i \boldsymbol{p}_i \,\middle|\, \phi_i \in \mathbb{R}\right\}$ its linear span and the subspace orthogonal to span(P) by $\operatorname{span}(P)^\perp = \left\{v \in \mathbb{R}^d \,\middle|\, \forall p \in \operatorname{span}(P) : \langle v, p\rangle = 0\right\}$;
- $\operatorname{aff}(P) = \left\{\sum_{i=1}^n \alpha_i \boldsymbol{p}_i \,\middle|\, \alpha_i \in \mathbb{R}, \sum_{i=1}^n \alpha_i = 1\right\}$ its affine hull;
- $\operatorname{pos}(P) = \left\{\sum_{i=1}^n \psi_i \boldsymbol{p}_i \,\middle|\, \psi_i \in \mathbb{R}_+\right\}$ all linear combinations with nonnegative coefficients. We call pos(P) the *positive span* of P and we call a combination with nonnegative coefficients a *positive combination*;
- $\operatorname{conv}(P) = \left\{\sum_{i=1}^n \lambda_i \boldsymbol{p}_i \,\middle|\, \lambda_i \in \mathbb{R}_+, \sum_{i=1}^n \lambda_i = 1\right\}$ its convex hull;
- $\dim P$ the dimension of span(P);
- and we call $\frac{1}{n}\sum_{i=1}^n \boldsymbol{p}_i$ the *barycenter* of P.

We say a set $Q \subseteq \mathbb{R}^d$ *embraces* a point $\boldsymbol{p} \in \mathbb{R}^d$ if $\boldsymbol{p} \in \operatorname{conv}(Q)$ and we say Q *ray-embraces* $\boldsymbol{p}$ if $\boldsymbol{p} \in \operatorname{pos}(Q)$. Furthermore, we denote with $\operatorname{int} Q = \left\{\boldsymbol{x} \in Q \mid \exists \varepsilon > 0 \text{ such that } B_\varepsilon(\boldsymbol{x}) \subseteq Q\right\}$ the *interior* of Q, where $B_\varepsilon(\boldsymbol{x}) \subset \mathbb{R}^d$ denotes the d-ball with radius ε that is centered at $\boldsymbol{x}$. Similarly, we denote with $\operatorname{relint} Q = \left\{\boldsymbol{x} \in Q \mid \exists \varepsilon > 0 \text{ such that } B_\varepsilon(\boldsymbol{x}) \cap \operatorname{aff}(Q) \subseteq Q\right\}$ the *relative interior* of Q.

If not otherwise noted, the algorithms in the first part of the thesis are analyzed in the WORD-RAM with logarithmic costs to comply with the definitions of the respective complexity classes (see [3, Theorem 1.4]). In the second part of the thesis, the algorithms are analyzed in the REAL-RAM [63, Section 1.4].

1.2. Descendants of the Colorful Carathéodory Theorem

Roughly 100 years ago, Carathéodory [16] proved a fundamental result about convex sets: every point in the convex hull of a point set $P \subset \mathbb{R}^d$ is also contained in a d-simplex with vertices in P. There are several elementary proofs and we repeat here a well-known constructive proof that leads to a polynomial-time algorithm for computing such a d-simplex.

Theorem 1.1 (Carathéodory's theorem)**.** *Let $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\} \subset \mathbb{R}^d$ be a set of n points.*

(Convex version) *If P embraces the origin, there is an affinely independent subset $P' \subseteq P$ that embraces the origin.*

(Cone version) *If P ray-embraces a point $\boldsymbol{b} \in \mathbb{R}^d$, there is a linearly independent subset $P' \subseteq P$ that ray-embraces $\boldsymbol{b}$.*

In the WORD-RAM, *the set P' can be computed in both versions in $O\big(\mathrm{poly}(d,n,L)\big)$ time, where L is the length of the input in binary. In the* REAL-RAM, *we assume to be given the coefficients of the convex combination of $\mathbf{0}$ with the points in P in case of the convex version, or the coefficients of the positive combination of $\boldsymbol{b}$ in case of the cone version. Then, we can compute P' in $O\big(d^3 n + n^2\big)$ time.*

Proof. The convex version can be easily reduced to the cone version as follows. Let $P \subset \mathbb{R}^{d-1}$ be a $\mathbf{0}$-embracing set and let $\widehat{P} \subset \mathbb{R}^d$ denote the point set that we obtain by appending a 1-coordinate to the points in P. Then, a linearly independent subset $\widehat{P}'$ of $\widehat{P}$ that ray-embraces the point $(0,\ldots,0,1)^T \in \mathbb{R}^d$ corresponds to an affinely independent subset P' of P that embraces the origin.

We prove the cone version by showing that if there is linear dependency among P, one point can be removed while preserving the property that $\boldsymbol{b}$ is ray-embraced. A repeated invocation of this argument implies the first part of the statement. First, if P is linearly independent, we have $|P| \le d$, so assume otherwise. Let $\boldsymbol{p}_1,\ldots,\boldsymbol{p}_n$ denote the points in P and let $\phi_1,\ldots,\phi_n \in \mathbb{R}$ be coefficients of a nontrivial linear combination of the origin, i.e.,

$$\mathbf{0} = \phi_1 \boldsymbol{p}_1 + \cdots + \phi_n \boldsymbol{p}_n \tag{1.1}$$

and not all ϕ_i, $i \in [n]$, are 0. We can assume without loss of generality that at least one ϕ_i is strictly greater than 0, since otherwise we can multiply all coefficients with -1. Furthermore, because $\boldsymbol{b} \in \mathrm{pos}(P)$, there are coefficients $\psi_1,\ldots,\psi_n \in \mathbb{R}_+$ such that

$$\boldsymbol{b} = \psi_1 \boldsymbol{p}_1 + \cdots + \psi_n \boldsymbol{p}_n. \tag{1.2}$$

Let $c \in \mathbb{R}$ be a factor that is to be specified. Scaling (1.1) by $c \in \mathbb{R}$ and subtracting it from (1.2), we obtain

$$\boldsymbol{b} = \sum_{i=1}^{n} \psi_i \boldsymbol{p}_i - c \sum_{i=1}^{n} \phi_i \boldsymbol{p}_i = \sum_{i=1}^{n} \psi'_i \boldsymbol{p}_i,$$

where $\psi'_i = \psi_i - c\phi_i$. Let $i^\star = \arg\min\{\psi_i/\phi_i \mid i \in [n], \phi_i > 0\}$, where ties are broken arbitrarily, and set $c = \psi_{i^\star}/\phi_{i^\star}$. Note that $i^\star$ is well defined since there exists at least one ϕ_i that is strictly greater than 0. Then, $\sum_{i=1}^n \psi'_i \boldsymbol{p}_i$ is a positive combination of $\boldsymbol{b}$ that involves at most $n-1$ points from P. Indeed by definition of $i^\star$, we have $\psi'_{i^\star} = (\psi_{i^\star} - c\phi_{i^\star}) = 0$ and for $i \neq i^\star$, we have $\psi'_i = (\psi_i - c\phi_i) \ge 0$.

It remains to show the running time. We begin with the WORD-RAM. In each iteration, we compute the coefficients ϕ_i and ψ_i, $i \in [n]$, with linear programming. Using the algorithm from [4], this takes $O\big(d^{1.5} n^{1.5} L\big)$ time, where L is the length of the input in binary. Then, the coefficients are encoded with $O\big(\mathrm{poly}(d,n,L)\big)$ bits, and hence finding the point $\boldsymbol{p}_{i^\star}$ needs $O\big(\mathrm{poly}(d,n,L)\big)$ time. Because there are $O(n)$ iterations, the total running time is $O\big(\mathrm{poly}(d,n,L)\big)$.

In the REAL-RAM, we compute in each iteration a linear dependency by Gaussian elimination in $O\big(d^3\big)$ time. By our assumption, we know the positive coefficients $\psi_1,\ldots,\psi_n$ and thus, we can find the point $\boldsymbol{p}_{i^\star} \in P$ in $O(n)$ time. Furthermore, we can compute the new coefficients $\psi'_i \in \mathbb{R}_+$, $i \in [n] \setminus \{i^\star\}$, from $\psi_1, \ldots, \psi_n$, the coefficients of the linear dependency, and the index $i^\star$ in $O(n)$ time. Hence, one iteration takes $O\left(d^3 + n\right)$ time and since there are $O(n)$ iterations, the algorithm needs in total $O\left(d^3 n + n^2\right)$ time. ■

Note that while the proof of Theorem 1.1 gives a polynomial-time algorithm, Knauer et al. [43] showed that deciding whether there exists a $\mathbf{0}$-embracing $(d-1)$-simplex is already NP-complete and W[1]-hard. In addition, it is known that for any $\varepsilon > 0$, there is a subset $Q \subseteq P$ of size only $O\left(1/\varepsilon^2\right)$ with $d_2(\mathbf{0}, \operatorname{conv}(Q)) = O(1/\varepsilon)$. Surprisingly, the constant in the $O(\cdot)$-notation does not depend on d. There is a particularly nice proof of this result by Blum et al. by a direct reduction to the analysis of the perceptron algorithm [13, Remark 2.8].

In 1982, Bárány [9] generalized Carathéodory's theorem by introducing colors: instead of only one set P that contains $\mathbf{0}$ in its convex hull, we now consider $d+1$ sets $C_1, \ldots, C_{d+1} \subset \mathbb{R}^d$ with $\mathbf{0} \in \operatorname{conv}(C_i)$ for $i = 1, \ldots, d+1$. We call the sets C_i *color classes* and we say a point $\boldsymbol{p}$ has color i if $\boldsymbol{p} \in C_i$. Then, we say a set $C \subseteq \bigcup_{i=1}^{d+1} C_i$ a *colorful choice* if it contains at most one point from each color class.[1] Bárány showed that there is always a $\mathbf{0}$-embracing colorful choice. This result is usually referred to as the *colorful Carathéodory theorem* [9, Theorem 2.1]. In the same publication, he also presented a further generalization: the *cone version* of the colorful Carathéodory theorem [9, Theorem 2.2]. To distinguish both versions, we refer to the first one as just the colorful Carathéodory theorem or the *convex version* of the colorful Carathéodory theorem. See Figure 1.1 for an example in two dimensions.

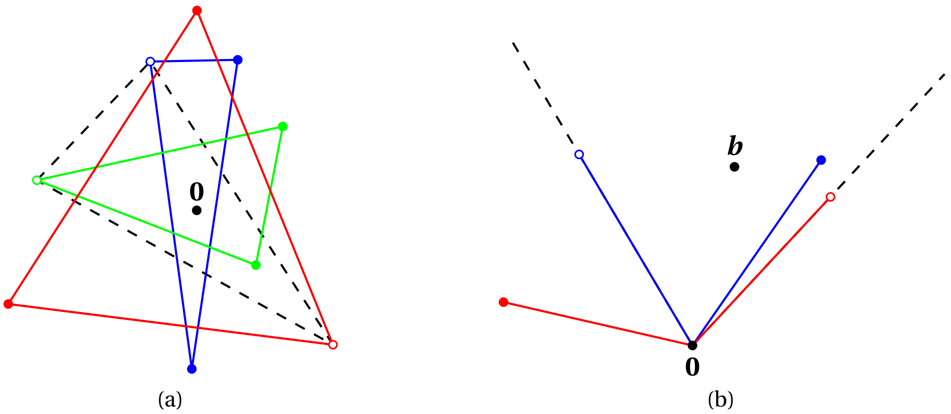

Figure 1.1.: (a) Example of the convex version of Theorem 1.2 in two dimensions. All color classes embrace the origin and the marked points form a $\mathbf{0}$-embracing colorful choice. (b) Example of the cone version of Theorem 1.2 in two dimensions. The two color classes ray-embrace $\boldsymbol{b}$ and the marked points form a colorful choice that ray-embraces $\boldsymbol{b}$, too.

Theorem 1.2 (Colorful Carathéodory theorem [9, Theorems 2.1 and 2.2])**.**

[1] If the color classes are not pairwise distinct, a colorful choice has to be defined more carefully. We say C is a colorful choice if it can be partitioned into $d+1$ sets C'_i, $i \in [d+1]$, such that $C'_i \subseteq C_i$ and $|C'_i| \leq 1$.

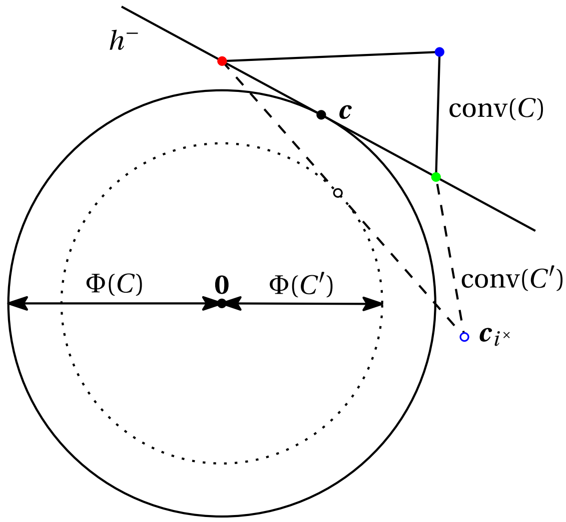

Figure 1.2.: Proof of the colorful Carathéodory theorem. If the potential function is larger than 0, it can always be decreased by swapping one point with another point of the same color.

(Convex version) *Let $C_1, \dots C_{d+1} \subset \mathbb{R}^d$ be $d+1$ finite sets that each embrace the origin. Then, there exists a $\mathbf{0}$-embracing colorful choice C.*

(Cone version) *Let $C_1, \dots, C_d \subset \mathbb{R}^d$ be d finite sets that ray-embrace a point $\boldsymbol{b} \neq \mathbf{0}$. Then, there exists a colorful choice C that ray-embraces $\boldsymbol{b}$.*

Proof. The convex version can be reduced to the cone version by using a similar lifting as described at the beginning of the proof of Theorem 1.1. For completeness, we sketch the proofs of both the convex version and the cone version as presented by Bárány.

We start with the convex version. Let C, $|C| \le d+1$, be a colorful choice of $C_1, \dots, C_{d+1}$. Let $\Phi(C)$ be the minimum ℓ_2-distance of a point in conv(C) to the origin. If $\Phi(C) = 0$, then $\mathbf{0} \in \text{conv}(C)$ and there is nothing left to show, so assume $\Phi(C) > 0$. Let $\boldsymbol{c}$ be the point in conv(C) with minimum ℓ_2-distance to the origin. Furthermore, let h^- be the halfspace that contains the origin and that is bounded by the hyperplane through $\boldsymbol{c}$ that is orthogonal to $\boldsymbol{c}$ interpreted as a vector. Since $\boldsymbol{c}$ minimizes the distance to the origin, it is contained in a facet of conv(C). Note that $\boldsymbol{c}$ is not necessarily contained in the interior of a facet. Then, Theorem 1.1 implies that there is a d-subset $F \subset C$ of C with $\boldsymbol{c} \in \text{conv}(F)$. Let $i^\times$ be the color of the point that is missing in F. The halfspace h^- contains the origin and thus it contains at least one point $\boldsymbol{c}_{i^\times} \in C_{i^\times}$ with color $i^\times$. Now, set $C' = (F \cup \{\boldsymbol{c}_{i^\times}\})$. Since it contains $\boldsymbol{c}$ and a point in h^-, we have $\Phi(C') < \Phi(C)$. Thus, the potential function Φ can always be strictly decreased if it is strictly larger than 0. The situation is depicted in Figure 1.2. Because there are only finitely many colorful choices, there exists then at least one colorful choice $C^\star$ with $\Phi(C^\star) = 0$.

We continue with the cone version. Again, let C, $|C| \le d$, denote some colorful choice. Here, let $\Phi(C)$ denote the minimum distance of a point in pos(C) to $\boldsymbol{b}$. Similar to the above argument, one can show that either $\Phi(C) = 0$ (in which case we are done) or that $\Phi(C)$ can be strictly decreased by swapping one point with another point of the same color. Again, since there is a finite number of colorful choices, this implies the statement. ■

Note that Carathéodory's theorem can be obtained directly from the colorful Carathéodory theorem by setting $C_1 = \cdots = C_{d+1} = P$. We define the corresponding computational problem as follows.

Definition 1.3 (COLORFULCARATHÉODORY)**.**

(Convex version)

GIVEN $d+1$ sets $C_1, \dots C_{d+1} \subset \mathbb{Q}^d$ that embrace the origin,

FIND a colorful choice that embraces the origin.

(Cone version)

GIVEN a point $\boldsymbol{b} \in \mathbb{Q}^d$, $\boldsymbol{b} \neq \mathbf{0}$, and d sets $C_1, \dots C_{d+1} \subset \mathbb{Q}^d$ that ray-embrace $\boldsymbol{b}$,

FIND a colorful choice that ray-embraces $\boldsymbol{b}$.

Algorithmic problems related to the colorful Carathéodory theorem have been first investigated by Bárány and Onn [11] in 1997. We discuss the results in more detail in Part II.

There are several generalizations of the colorful Carathéodory theorem. Independently, Arocha et al. [5, Theorem 1] and Holmsen et al. [37, Theorem 8] showed that it is enough if the origin is contained in the convex hull of each pair of color classes C_i, C_j for $1 \leq i < j \leq d+1$ to guarantee the existence of a $\mathbf{0}$-embracing colorful choice. This was further generalized by Meunier and Deza [53, Theorem 3] who showed that a $\mathbf{0}$-embracing colorful choice already exists if for each pair of color classes C_i, C_j, $1 \leq i < j \leq d+1$, there exists a color class C_k, $k \neq i, j$, such that for all $\boldsymbol{p} \in C_k$ the ray originating at $\boldsymbol{p}$ through the origin intersects $\operatorname{conv}(C_i \cup C_j)$ in a point different from $\boldsymbol{p}$. Moreover, they showed that an ever weaker assumption suffices. Given a d-set $P \subset \mathbb{R}^d$, we denote with $h^-(P)$ the halfspace that contains $\mathbf{0}$ and is bounded by $\operatorname{aff}(P)$. Furthermore, we call a set $T_{\bar{i}}$ that contains exactly one point from color classes $C_1, \dots, C_{i-1}, C_{i+1}, \dots, C_{d+1}$ and *no* point from color class C_i an *$\bar{i}$-traversal.* Now, a $\mathbf{0}$-embracing colorful choice exists if for all pairs of color classes C_i, C_j, $1 \leq i < j \leq d+1$, and all $\bar{i}$-traversals $T_{\bar{i}}$ at least one point of $C_i \cup C_j$ is contained in $h^-(T_{\bar{i}})$ [53, Theorem 4]. These generalizations form in fact a strict hierarchy in the order as presented here [53]. Note that the preconditions of the generalizations except for the last one are polynomial-time checkable.

Theorem 1.2 itself only guarantees the existence of at least *one* $\mathbf{0}$-embracing colorful choice. Let us denote by $\mu(d)$ the minimum number of $\mathbf{0}$-embracing colorful choices over all configurations of color classes $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ such that for all $i \in [d+1]$, the set C_i has size $d+1$ and embraces the origin, and such that the set $\{\mathbf{0}\} \cup \bigcup_{i=1}^{d+1} C_i$ is in general position (i.e., no $k+2$ points lie in a common k-flat for $k \in [d-1]$). Bárány already showed that each point $\boldsymbol{p} \in C_1 \cup \cdots \cup C_{d+1}$ is contained in some colorful $\mathbf{0}$-embracing colorful choice [9, Theorem 2.3]. This directly implies $\mu(d) \geq d+1$. In 2006, Deza et al. improved the lower bound to $\mu(d) \geq 2d$ [29, Corollary 3.10] and gave examples where the origin is contained in only d^2+1 colorful choices that embrace the origin [29, Section 3.4]. They conjectured d^2+1 to be the true value of $\mu(d)$, which became known as the *colorful simplicial depth conjecture*. One year later, Bárány and Matoušek [10, Theorem 1.2] proved this conjecture asymptotically by showing that $\mu(d) \geq \frac{1}{5}d(d+1)$. In 2014, Deza et al. [30, Proposition 5.2] proved the conjecture for $d = 4$ and a year later, Sarrabezolles [71, Theorem 1] proved the conjecture for all $d \in \mathbb{N}$.

A result that can be seen as a dual to the cone version of the colorful Carathéodory theorem is the colorful Helly theorem due to Lovász [47]. First, we shortly restate the *non-colorful* Helly theorem [36]. Let $\mathcal{C}$ be a finite family of compact convex sets in $\mathbb{R}^d$ such that every subfamily $\mathcal{C}' \subseteq \mathcal{C}$ of size $d+1$ has a nonempty intersection, then the complete family $\mathcal{C}$ has a nonempty intersection. Now, Lovász proved the following colorful generalization.

Theorem 1.4 (Colorful Helly theorem [47])**.** *Let $\mathcal{C}_1, \ldots, \mathcal{C}_{d+1}$ be finite families of compact convex sets in $\mathbb{R}^d$ such that for all colorful choices $\mathcal{C} = \{K_1, \ldots, K_{d+1}\}$, where $K_i \in \mathcal{C}_i$ for $i \in [d+1]$, we have $\bigcap_{i=1}^{d+1} K_i \neq \emptyset$. Then, there is a set family $\mathcal{C}_i$ whose sets have a nonempty intersection, where $i \in [d+1]$.*

We prove Theorem 1.4 by a reduction to the cone version of the colorful Carathéodory theorem, where we follow Bárány's presentation. The following two lemmas are the key component to relate both theorems. The first one enables us to approximate convex sets by halfspaces.

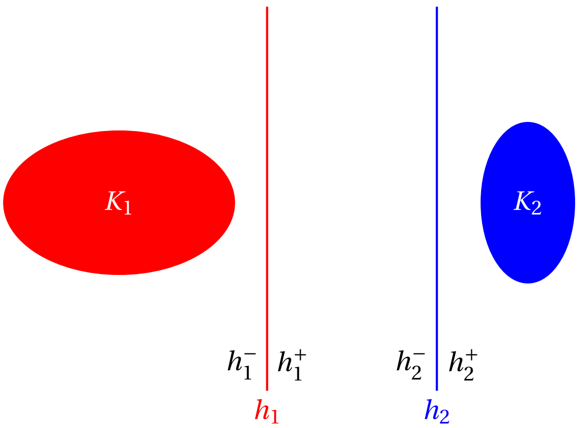

Figure 1.3.: Proof situation of Lemma 1.5 for $n = 2$.

Lemma 1.5 ([Folklore])**.** *Let $K_1, \ldots, K_n \subset \mathbb{R}^d$, $n \geq 2$, be compact convex sets with $\bigcap_{i=1}^{n} K_i = \emptyset$. Then, there exist halfspaces $h_1^-, \ldots, h_n^-$ such that $K_i \subseteq h_i^-$ for $i \in [n]$ and $\bigcap_{i=1}^{n} h_i^- = \emptyset$.*

Proof. We first consider the case $n = 2$. Since both sets K_1, K_2 are compact convex sets, the separation theorem [48, Theorem 1.2.4] guarantees the existence of a strictly separating hyperplane h_1 such that $K_1 \subset h_1^-$ and $K_2 \subset h_1^+$, where h_1^- and h_1^+ denote the two halfspaces that are bounded by h_1. Furthermore, since h_1 and K_2 are convex sets and since K_2 is compact, there exists a strictly separating hyperplane h_2 such that $h_1 \subset h_2^-$ and $K_2 \subset h_2^+$, where h_2^- and h_2^+ denote the two halfspaces that are bounded by h_2. Note that h_1 and h_2 must be parallel. Hence, $h_1^- \cap h_2^+ = \emptyset$ and $K_1 \subset h_1^-$, $K_2 \subset h_2^+$ as desired. Please see Figure 1.3.

Now let $n > 2$. We iteratively construct halfspaces $h_1^-, \dots, h_j^-$, $j \in [n]_0$ while maintaining the invariant that

$$K_i \subseteq h_i^- \text{ for } i \in [j] \text{ and } \left(\bigcap_{i=1}^{j} h_i^-\right) \cap \left(\bigcap_{i=j+1}^{n} K_i\right) = \emptyset.$$

Assume we have already constructed $h_1^-, \dots, h_j^-$ for some fixed $j \in [n-1]_0$ and we want to construct the halfspace h_{j+1} such that $K_{j+1} \subseteq h_{j+1}^-$ and $\left(\bigcap_{i=1}^{j} h_i^-\right) \cap \left(\bigcap_{i=j+1}^{n} K_i\right) = \emptyset$. Let $K_{j+1}^\times$ denote the intersection

$$K_{j+1}^\times = \left(\bigcap_{i=1}^{j} h_i^-\right) \cap \left(\bigcap_{i=j+2}^{n} K_i\right)$$

that misses K_{j+1}. Since $\left(\bigcap_{i=1}^{j} h_i^-\right) \cap \left(\bigcap_{i=j+1}^{n} K_i\right) = \emptyset$, we have $K_{j+1} \cap K_{j+1}^\times = \emptyset$. By the above argument, there then exists a halfspace h_{j+1}^- such that $K_{j+1} \subseteq h_{j+1}^-$ and

$$\emptyset = h_{j+1}^- \cap K_{j+1}^\times = \left(\bigcap_{i=1}^{j+1} h_i^-\right) \cap \left(\bigcap_{i=j+2}^{n} K_i\right).$$

Hence, there exist halfspaces $h_1^-, \dots, h_n^-$ such that $K_i \subseteq h_i^-$ for $i \in [n]$ and $\bigcap_{i=1}^{n} h_i^- = \emptyset$, as claimed. ■

The next lemma shows that if a set of halfspaces has an empty intersection, then the dual point set contains a specific point in its positive span.

Lemma 1.6 ([Folklore])**.** *Let $\mathcal{H} = \{h_1^-, \dots, h_n^-\}$ be a set of halfspaces, where each halfspace $h_i^- \subset \mathbb{R}^d$ is defined as $h_i^- = \{\boldsymbol{x} \in \mathbb{R}^d \mid \boldsymbol{a}_i \boldsymbol{x} \leq \alpha_i\}$ with $\boldsymbol{a}_i \in \mathbb{R}^d$ and $\alpha_i \in \mathbb{R}$. Let $\overline{\mathcal{H}}$ denote the set of points $\{\overline{\boldsymbol{a}_i} = (\boldsymbol{a}_i, \alpha_i)^T \in \mathbb{R}^{d+1} \mid h_i^- \in \mathcal{H}\}$ dual to $\mathcal{H}$. Then, the point $\boldsymbol{b} = (0, \dots, 0, -1)^T$ is contained in $\operatorname{pos}\left(\overline{\mathcal{H}}\right)$ if and only if the intersection of halfspaces $\bigcap_{i=1}^{n} h_i^-$ is empty.*

Proof. Assume that $\boldsymbol{b} \in \operatorname{pos}(\overline{\mathcal{H}})$. Then, there exist coefficients $\psi_1, \dots, \psi_n \in \mathbb{R}^+$, not all 0, such that $\boldsymbol{b} = \sum_{i=1}^{n} \psi_i \overline{\boldsymbol{a}_i}$. In particular, we have $\sum_{i=1}^{n} \psi_i \boldsymbol{a}_i = \boldsymbol{0}$ and $\sum_{i=1}^{n} \psi_i \alpha_i = -1$. This directly implies that the following system of linear inequalities

$$\begin{gathered} \boldsymbol{a}_1 \boldsymbol{x} \leq \alpha_1 \\ \vdots \\ \boldsymbol{a}_n \boldsymbol{x} \leq \alpha_n \end{gathered}$$

is inconsistent and hence the intersection of the halfspaces $h_1^-, \dots, h_n^-$ is empty.

Now assume that $\boldsymbol{b} \notin \operatorname{pos}(\overline{\mathcal{H}})$. We show that this implies a nonempty intersection $\bigcap_{i=1}^{n} h_i^-$ of the halfspaces. By our assumption, there exists a strictly separating hyperplane

$$h = \left\{\boldsymbol{x} \in \mathbb{R}^{d+1} \,\middle|\, \boldsymbol{y}^T \boldsymbol{x} = \gamma\right\}, \boldsymbol{y} \in \mathbb{R}^{d+1}, \gamma \in \mathbb{R}_+,$$

with $\boldsymbol{y}^T \boldsymbol{b} > \gamma$ and $\boldsymbol{y}^T \overline{\boldsymbol{a}_i} < \gamma$ for $i \in [n]$. For $\lambda \in [0,1]$, let $h_\lambda \subset \mathbb{R}^{d+1}$ denote the hyperplane $\{\boldsymbol{x} \in \mathbb{R}^{d+1} \mid \boldsymbol{y}^T \boldsymbol{x} = \lambda\gamma\}$. Choose $\lambda \in [0,1]$ maximum such that h_λ intersects $\operatorname{pos}(\overline{\mathcal{H}})$ and assume for the sake of contradiction that $\lambda > 0$. Then, we have $\boldsymbol{y}^T\overline{\boldsymbol{a}_i} = \lambda\gamma$ for some $i \in [n]$. Thus, h contains the point $\lambda^{-1}\overline{\boldsymbol{a}_i} \in \operatorname{pos}(\overline{\mathcal{H}})$, a contradiction to h strictly separating $\boldsymbol{b}$ and $\operatorname{pos}(\overline{\mathcal{H}})$. Hence, the hyperplane $h_\star = \left\{\boldsymbol{x} \in \mathbb{R}^{d+1} \,\middle|\, \boldsymbol{y}^T\boldsymbol{x} = 0\right\}$ contains the origin and separates $\boldsymbol{b}$ and $\operatorname{pos}(\overline{\mathcal{H}})$.

Let now $\boldsymbol{y}' \in \mathbb{R}^d$ denote the first d coordinates of $\boldsymbol{y}$ and let γ' denote the $(d+1)$th coordinate of $\boldsymbol{y}$. Because $\gamma \geq 0$, we have $\boldsymbol{y}^T\boldsymbol{b} > 0$ and hence $\gamma' < 0$. Since $\boldsymbol{y}^T\overline{\boldsymbol{a}_i} \leq 0$, we have

$$\left(\boldsymbol{y}'\right)^T \boldsymbol{a}_i \leq -\gamma'\alpha_i = \left|\gamma'\right|\alpha_i \text{ for } i \in [n].$$

Hence, that the point $\boldsymbol{y}'/\left|\gamma'\right|$ is contained in the intersection of all halfspaces $h_1^-,\ldots,h_n^-$. ■

Equipped with Lemmas 1.5 and 1.6, we are now ready to prove the colorful Helly theorem.

Proof of Theorem 1.4. We show the contrapositive of the statement, i.e., if the sets of each color class $\mathcal{C}_i$, $i \in [d+1]$, have an empty intersection, then there exists a colorful choice $\mathcal{C} = \{K_1,\ldots,K_{d+1}\}$ with $\bigcap_{i=1}^{d+1} K_i = \emptyset$.

In a first step, we replace each color class $\mathcal{C}_i$, $i \in [d+1]$, by at most $d+1$ halfspaces as follows. Helly's theorem [48, Theorem 1.3.2] states we can find a subset $\mathcal{C}_i' = \{K_{i,1},\ldots,K_{i,n_i}\} \subseteq \mathcal{C}_i$, where $n_i \leq d+1$, whose sets have an empty intersection. By applying Lemma 1.5, we obtain a set of n_i halfspaces $\mathcal{H}_i = \{h_{i,1}^-,\ldots,h_{i,n_i}^-\}$ with an empty intersection and each halfspace $h_{i,j}^-$ contains its corresponding convex set $K_{i,j}$. Furthermore, we color the halfspaces in $\mathcal{H}_i$ with color i for $i \in [d+1]$. Now, for each set of halfspaces $\mathcal{H}_i$, $i \in [d+1]$, let $C_i = \overline{\mathcal{H}_i} \subset \mathbb{R}^{d+1}$ denote the set of points dual to $\mathcal{H}_i$ as in Lemma 1.6. Since each set $\mathcal{H}_i$, $i \in [d+1]$, has an empty intersection, the dual set C_i ray-embraces $\boldsymbol{b} = (0,\ldots,0,-1)^T \in \mathbb{R}^{d+1}$. We color the points in C_i, $i \in [d+1]$, with color i. Then, the cone version of Theorem 1.2 guarantees the existence of a colorful choice $C \subset \bigcup_{i=1}^{d+1} C_i$ that ray-embraces $\boldsymbol{b}$. Again by Lemma 1.6, the colorful choice C of dual points corresponds to a colorful choice $\mathcal{H}_C \subset \bigcup_{i=1}^{d+1} \mathcal{H}_i$ of halfspaces with $\overline{\mathcal{H}_C} = C$ that has an empty intersection. Then, $\mathcal{H}_C$ corresponds in turn to a colorful choice $\mathcal{C} \subset \bigcup_{i=1}^{d+1} \mathcal{C}_i'$ of convex sets that has an empty intersection, which concludes the proof. ■

Note that Helly's theorem can be obtained from the cone version of Carathéodory's theorem in the same fashion as in the proof of Theorem 1.4.

A surprising application of the colorful Carathéodory theorem appears in the context of Tukey depth. Rado [65] showed in his well-known centerpoint theorem that there always exist a point with Tukey depth linear in the size of P, a centerpoint.

Theorem 1.7 (Centerpoint theorem [65, 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 $\tau \geq \left\lceil \frac{|P|}{d+1} \right\rceil$.* ■

It can be easily seen that this bound is tight: consider a regular simplex $\sigma \subset \mathbb{R}^d$ and replace each vertex of the simplex by a "small" point cloud of size $\frac{n}{d+1}$, where $n \in \mathbb{N}$ is some multiple of $d+1$. Then any point $\boldsymbol{q}$ in σ is contained in a halfspace that only contains the point cloud of one vertex.

Teng [74, Theorem 8.4] showed that given a point set $P \in \mathbb{R}^d$ and a candidate centerpoint $\boldsymbol{q} \in \mathbb{R}^d$ it is coNP-complete to decide whether $\boldsymbol{q}$ is a centerpoint of P if d is not constant. For

$d = 1$, a centerpoint is equivalent to a median of a set of numbers and hence can be computed in $O(|P|)$ time [14]. Jadhav and Mukhopadhyay [38] showed that linear time is even sufficient in two dimensions. For $d \geq 3$ fixed, the best known algorithm is by Chan [17] who showed how to compute a point with maximum Tukey depth in expected time $O\left(n^{d-1}\right)$.

As stated at the beginning, Tverberg partitions serve as a polynomial-time checkable certificate for a subset of centerpoints: Tverberg points. In recent years, this property has been exploited algorithmically to derive efficient approximation algorithms for centerpoints [57, 59]. The existence of Tverberg points is guaranteed by Tverberg's theorem [76].

Theorem 1.8 (Tverberg's theorem [76])**.** *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.*

While Tverberg's first proof is quite involved, several simplified subsequent proofs [67, 69, 77, 78] have been published. Here, we present Sarkaria's proof [69] with further simplifications from Bárány and Onn [11], and Arocha et al. [5]. The main tool is the next 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. [5]. 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

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

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

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

Lemma 1.9 (Sarkaria's lemma [69], [5, Lemma 2])**.** *Let $P_1, \ldots, P_m \subset \mathbb{R}^d$ be m point sets and let $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_m \subset \mathbb{R}^{m-1}$ be m vectors with $\boldsymbol{q}_i = \boldsymbol{e}_i$ for $i \in [m-1]$ and $\boldsymbol{q}_m = -\mathbf{1}$. 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 $\hat{\boldsymbol{p}}_i \in$

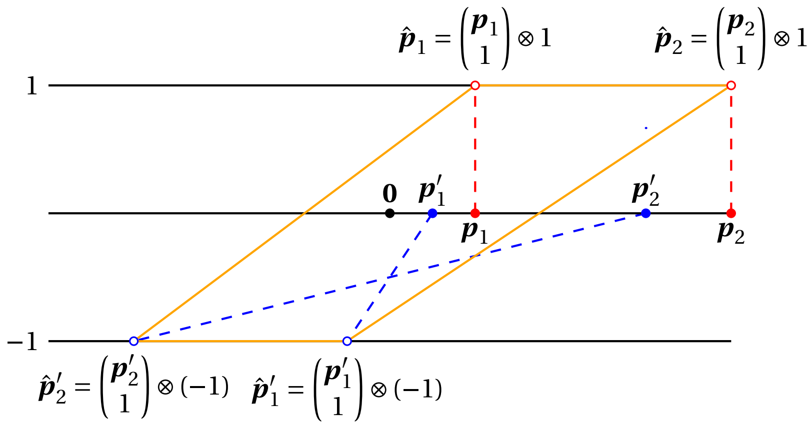

Figure 1.4.: An example of Sarkaria's lemma for $d = 1$ and $m = 2$. The set P_1 consists of the red points and the set P_2 consists of the blue points. Since the convex hulls of P_1 and P_2 intersect, the lifted points embrace the origin.

$\operatorname{conv}\big(\widehat{P}_i\big)$, $i \in [m]$, that we obtain by using the same convex coefficients for the points in $\widehat{P}_i$, i.e., set

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

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

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

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

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

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

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

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 $\operatorname{conv}(P_1), \ldots, \operatorname{conv}(P_m)$. ■

Please refer to Figure 1.4 for an example of Sarkaria's lifting argument. Little work is now left to obtain Tverberg's theorem from Lemma 1.9 and the colorful Carathéodory theorem.

Proof of Theorem 1.8. Let $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\} \subset \mathbb{R}^d$ be a point set of size $n = (d+1)(m-1)+1$ and let $P_1, \ldots, P_m$ 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 1.9, 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 1.9 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 1.2, 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 1.9 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. ■

Interestingly, to prove the centerpoint theorem already Helly's theorem (and thus the cone version of Carathéodory's theorem) is sufficient.

Even less effort is required to obtain the colorful Kirchberger theorem from Lemma 1.9. Let $A, B \subset \mathbb{R}^d$ be two point sets. Kirchberger's theorem [42] states that if for all subsets $C \subset A \cup B$ of size at most $d+2$, the sets $\operatorname{conv}(A \cap C)$ and $\operatorname{conv}(B \cap C)$ have an empty intersection, then $\operatorname{conv}(A)$ and $\operatorname{conv}(B)$ have an empty intersection. It can be proven using a Helly-type theorem [64, Section 2] or a Carathéodory-type theorem [81]. Arocha et al. [5] presented a generalization based on the colorful Carathéodory theorem.[2]

Theorem 1.10 (Colorful Kirchberger theorem [5, special case of Theorem 3])**.** *Let $C_1, \ldots, C_n \subset \mathbb{R}^d$ be $n = (m-1)(d+1)+1$ pairwise disjoint color classes and let $\mathcal{T}_i = \{T_{i,1}, \ldots, T_{i,m}\}$ denote a Tverberg m-partition for C_i, where $i \in [n]$. Then, there exists a colorful choice C, $|C| = n$, such*

[2] Actually, Arocha et al. present an even stronger result (the "very colorful Kirchberger theorem" [5, Theorem 3]) using a generalization of the colorful Carathéodory theorem. Here, we consider the weaker version that can be obtained from Theorem 1.2.

that the family of sets

$$\mathcal{T}_C = \left\{ C \cap \left(\bigcup_{i=1}^{n} T_{i,j} \right) \,\middle|\, j \in [m] \right\}$$

is a Tverberg m-partition for C.

Proof. We lift each Tverberg partition to $\mathbb{R}^{n-1}$ as in Lemma 1.9: for $i \in [n]$ and $j \in [m]$, we denote with $\widehat{T}_{i,j}$ the set

$$\widehat{T}_{i,j} = \left\{ \boldsymbol{p} \otimes \boldsymbol{q}_j \,\middle|\, \boldsymbol{p} \in T_{i,j} \right\} \subset \mathbb{R}^{n-1}.$$

By Lemma 1.9 and since each set $\mathcal{T}_i$, $i \in [n]$, is a Tverberg partition, the sets $\widehat{C}_i = \bigcup_{j=1}^{m} \widehat{T}_{i,j}$, $i \in [n]$, capture the origin. We color the points in $\widehat{C}_i$ with color i. Since there are n color classes that capture the origin in $n-1$ dimensions, Theorem 1.2 guarantees the existence of a colorful choice $\widehat{C}$ that embraces the origin. For $j \in [m]$, let $\widehat{T}_j = \widehat{C} \cap \left(\bigcup_{i=1}^{n} \widehat{T}_{i,j} \right)$ denote all points from a jth element in a Tverberg partition in C. Since $\widehat{C} = \bigcup_{j=1}^{m} \widehat{T}_j$ embraces the origin, Lemma 1.9 implies that the convex hulls of the sets $T_j = \left\{ \boldsymbol{p} \in \bigcup_{i=1}^{n} P_i \,\middle|\, \boldsymbol{p} \otimes \boldsymbol{q}_j \in \widehat{T}_j \right\}$ have a nonempty intersection. Further, since for $j \in [m]$, the set $\widehat{T}_j$ is a subset of $\bigcup_{i=1}^{n} \widehat{T}_{i,j}$, we have $T_j \subset \left(\bigcup_{i=1}^{n} T_{i,j} \right)$. Moreover, since all points that correspond to the Tverberg partition $\mathcal{T}_i$, $i \in [n]$, have color i, exactly one of the sets $T_1, \ldots, T_m$ contains a point from C_i. ■

Similar to the Tukey depth, the simplicial depth [46] is a notion of data depth. Given a set of points $P \subset \mathbb{R}^d$, the *simplicial depth* $\sigma_P(\boldsymbol{q}) \in \mathbb{N}$ of a point $\boldsymbol{q} \in \mathbb{R}^d$ with respect to the set P is the number of d-simplices with vertices in P that contain $\boldsymbol{q}$. More formally, the simplicial depth is defined as

$$\sigma_P(\boldsymbol{q}) = \left| \left\{ S \subseteq P \,\middle|\, \boldsymbol{q} \in \operatorname{conv}(S), |S| = d+1 \right\} \right|.$$

Let $\sigma(n,d) = \min_{P \subset \mathbb{R}^d, |P|=n} \max_{q \in \mathbb{R}^d} \sigma_P(q)$ denote the minimum simplicial depth that can always be achieved by some point for point sets of size n in d dimensions. The first selection lemma guarantees that there is always a point with simplicial depth $\Omega\left(|P|^{d+1}\right)$ if d is fixed. This is asymptotically tight as there are only $O\left(|P|^{d+1}\right)$ candidate simplices in total. In fact, as we discuss below, centerpoints maximize asymptotically the simplicial depth. The first selection lemma for arbitrary but fixed d was first shown by Bárány [9, Theorem 5.1]. There are several ways to prove this lemma: Bárány's proof employs Tverberg's theorem and the colorful Carathéodory theorem to show that every Tverberg point has simplicial depth $\Omega\left(|P|^{d+1}\right)$. A different proof [48, Theorem 9.1.1 (second proof)] is based on Tverberg's theorem and the fractional Helly theorem, however it only shows existence of such a point and does not relate it with the Tukey depth.

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

Lemma 1.11. *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 1.2, 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. ■

Note that Lemma 1.11 also shows that a point $\boldsymbol{c}'$ with Tukey depth τ has simplicial depth $\Omega\left(\frac{\tau^{d+1}}{d^{2d+2}}\right)$. This can be seen as follows: by repeatedly applying Theorem 1.1, we can build a Tverberg partition $\mathcal{T}$ of size $\lfloor \tau/d \rfloor$ such that $\boldsymbol{c}' \in \bigcap_{T\in\mathcal{T}} \operatorname{conv}(T)$. Then, Lemma 1.11 guarantees that $\boldsymbol{c}'$ has the claimed simplicial depth. This observation is folklore and its algorithmic implications have been recently explored by Rolnick and Sobéron [66].

The first selection lemma is now an immediate consequence of Lemma 1.11 and Theorem 1.8, or of Lemma 1.11 and Theorem 1.7.

Theorem 1.12 (First selection lemma [9, 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 $\sigma_P(\boldsymbol{q}) = \Omega\left(|P|^{d+1}\right)$.*

The constant c_d in the $\Omega(\cdot)$-notation is not yet settled. Wagner [80, Section 4.4] showed that for centerpoints, c_d is lower bounded by $\frac{d^2+1}{(d+1)!(d+1)^{d+1}}$. Bukh, Matoušek, and Nivasch [15, Theorem 1.3] proved that $c_d \le 1/(d+1)^{d+1}$ and they conjectured this bound to be tight. In particular, it coincides with the known value for $c_2 = 1/27$ in the plane. See Figure 1.5 for an overview of the presented proofs.

We conclude the introduction by defining the computational problems that correspond to the presented implications of the colorful Carathéodory theorem.

Definition 1.13. 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.

- COLORFULKIRCHBERGER

  GIVEN $n = (m-1)(d+1)+1$ pairwise disjoint color classes $C_1, \ldots, C_n \subset \mathbb{Q}^d$, each of size n, and for each color class C_i, $i \in [n]$, a Tverberg m-partition $\mathcal{T}_i = \{T_{i,1}, \ldots, T_{i,m}\}$,

  FIND a colorful choice C, $|C| = n$, such that the family of sets

  $$\mathcal{T}_C = \left\{ C \cap \left(\bigcup_{i=1}^{n} T_{i,j} \right) \middle| \, j \in [m] \right\}$$

  is a Tverberg m-partition for C.

- SIMPLICIALCENTER

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

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

Interpreting the presented proofs of the centerpoint theorem, Tverberg's theorem, the colorful Kirchberger theorem, and the first selection lemma as algorithms, we obtain the following lemma.

Lemma 1.14. *In the following, let L denote the length of the input in binary. In the* WORD-RAM, *given access to an oracle for* COLORFULCARATHÉODORY, TVERBERG *can be solved in $O(m^2 L)$ additional time. Furthermore,* COLORFULKIRCHBERGER *can be solved in $O(mL)$ additional time, and both* CENTERPOINT *and* SIMPLICIALCENTER *can be solved in $O(n^3 L)$ additional time.*

Proof. As shown in the proof of Theorem 1.8, to compute a Tverberg partition, it suffices to lift $m = \lceil \frac{n}{d+1} \rceil$ copies of the input point set $P \subset \mathbb{Q}^d$ with Lemma 1.9 and then query the oracle for COLORFULCARATHÉODORY. Lifting one point set needs $O(mL)$ time and hence we need $O(m^2 L)$ 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 [4], we can compute a Tverberg point in $O(n^3 L)$ time by solving the linear system

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

where L is the length of the input in binary.

To solve COLORFULKIRCHBERGER, we need to lift all points in $\bigcup_{i=1}^{n} C_i$ with Lemma 1.9 and then query the oracle for COLORFULCARATHÉODORY, as presented in the proof of Theorem 1.10. This needs $O(mL)$ time in total. ■

From now on, we refer to CENTERPOINT, TVERBERG, COLORFULKIRCHBERGER, and SIMPLICIALCENTER as the *descendants* of COLORFULCARATHÉODORY. See Figure 1.6 for an overview of the reductions.

1.3. Previous Work

Chapters 3, 4, and 8, as well as Section 5.1 are based on joint work with Frédéric Meunier, Wolfgang Mulzer, and Pauline Sarrabezolles and have been published in [54].

Chapter 7, Chapter 9, and Section 5.2 are based on joint work with Wolfgang Mulzer and have been published in [58].

Chapter 10 is based on joint work with Peyman Afshani and Donald R. Sheehy and is based on the technical report [2].

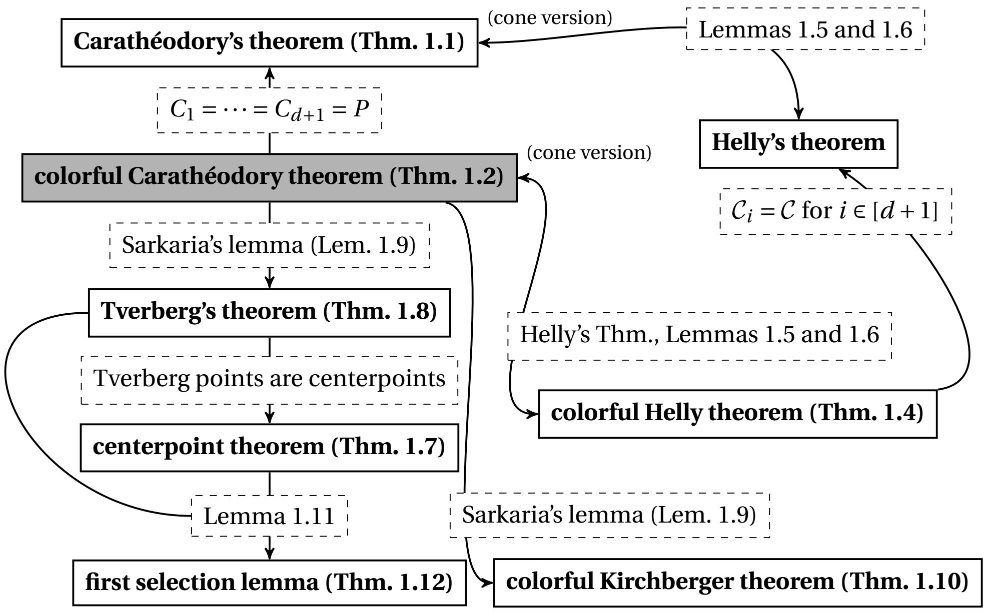

Figure 1.5.: Implications of the colorful Carathéodory theorem. Arrows represent implications.

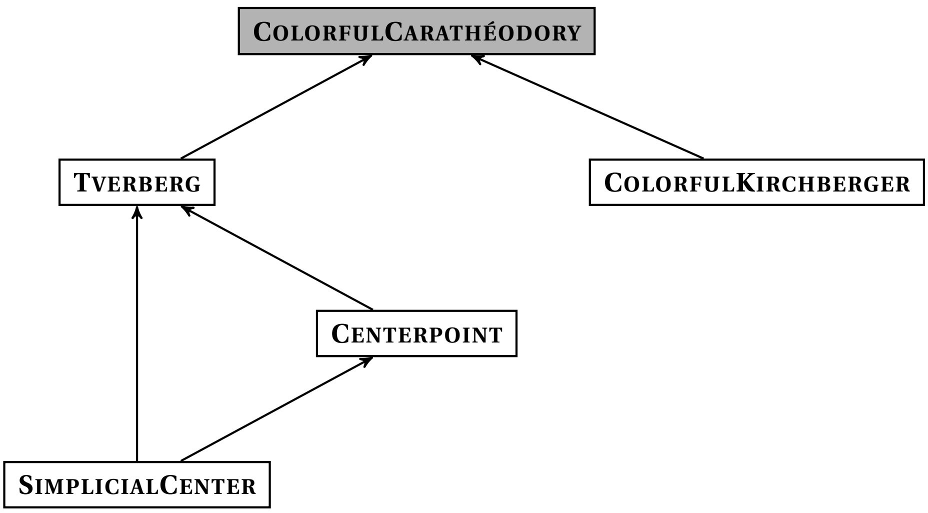

Figure 1.6.: Web of reductions. Arrows represent polynomial-time reductions.

I

The Complexity of COLORFULCARATHÉODORY

2 Introduction to Part I

The complexity of ColorfulCarathéodory and its descendants is still widely open apart from a trivial upper bound on the complexity, the class *total function NP*. Our ultimate goal in this part of the thesis is to improve this upper bound on the complexity of Colorful-Carathéodory and thereby to derive new upper bounds on the complexity of its descendants.

In all considered problems, we can check the preconditions that guarantee the existence of a solution in polynomial time, however we know of no polynomial-time algorithms that can actually find solutions. The classic way to model the computational complexity of a problem is to study its corresponding decision problem: *does there exist a solution?* Or in the case of optimization problems: *does there exist a solution of a specific quality*? However, in the case of ColorfulCarathéodory, the corresponding decision problem is trivial. This led to the study of related decision problems. Meunier and Sarrabezolles [55, Theorem 2] showed that the problem

Given k color classes $C_1, \ldots, C_k \subset \mathbb{Q}^d$ that do not necessarily embrace the origin,

Decide whether there exists a $\mathbf{0}$-embracing colorful choice,

is NP-complete if d is part of the input. We discuss an alternative proof for this result in Section 5.2. Related to Tverberg, Teng [74, Theorem 8.8] showed that the decision problem

Given a point set $P \subset \mathbb{Q}^d$, and a point $\boldsymbol{q} \in \mathbb{Q}^d$

Decide whether $\boldsymbol{q}$ is a Tverberg point for P,

is NP-complete. Finally, as discussed before, Teng [74, Theorem 8.4] showed that the problem

Given a point set $P \subset \mathbb{Q}^d$, and a point $\boldsymbol{q} \in \mathbb{Q}^d$

Decide whether $\boldsymbol{q}$ is a centerpoint of P.

is coNP-complete if d is part of the input. Although related, hardness of these decision problems do not necessarily tell us anything about the complexity of computing *any* solution to ColorfulCarathéodory or its descendants. A class designed to capture the complexity of search problems is *function NP* (FNP), an analogue of NP that contains search problems for which we can verify a solution in polynomial time. More formally, FNP consists of binary relations $\mathcal{R}$ between a set of problem instances $\mathcal{I} \subseteq \{0,1\}^\star$ and a set of candidate solutions $\mathcal{S} \subseteq \{0,1\}^\star$ such that

- the length of the candidate solutions is polynomially bounded in the length of the instances, i.e., there is a polynomial p such that for all pairs $(I, s) \in \mathcal{R}$, we have

$$|s| \leq p(|I|);$$

- $\mathcal{I}$ is polynomial-time verifiable;
- the relation $\mathcal{R}$ is polynomial-time verifiable.

The computational problem that corresponds to a relation $\mathcal{R} \in \mathsf{FNP}$ is given an instance $I \in \mathcal{I}$, find a solution $s \in \mathcal{S}$ such that $(I, s) \in \mathcal{R}$ if it exists and otherwise state that there is none. We call a relation $\mathcal{R} \in \mathsf{FNP}$ *total* if for all $I \in \mathcal{I}$ there exists some $s \in \mathcal{S}$ with $(I, s) \in \mathcal{R}$. The subclass *total function NP* (TFNP) [52] consists of all total relations in FNP. This class can be seen as a search-problem analogue to $\mathsf{NP} \cap \mathsf{coNP}$ (as described in [52, Section 1.1]). FP, the search-problem analogue to P and a subclass of TFNP, consists of all relations $\mathcal{R}$ in FNP for which there exists a deterministic polynomial-time Turing machine that, on input $I \in \mathcal{I}$, outputs some $s \in \mathcal{S}$ with $(I, s) \in \mathcal{R}$.

Now, given a problem instance of the cone version of ColorfulCarathéodory, i.e., given d color classes $C_1, \dots, C_d \subset \mathbb{Q}^d$ and a vector $\boldsymbol{b} \in \mathbb{Q}^d$, we can verify in polynomial time whether each color class ray-embraces $\boldsymbol{b}$ using linear programming. Furthermore, given a colorful choice C, we can verify in polynomial time whether C ray-embraces $\boldsymbol{b}$. Hence, Colorful-Carathéodory is in TFNP. Then, Lemma 1.14 immediately puts Centerpoint, Tverberg, ColorfulKirchberger, and SimplicialCenter in TFNP. The following theorem shows that this has already a nontrivial consequence.

Theorem 2.1 ([52, Theorem 2.1], [39, Lemma 4])**.** *No problem in TFNP is NP-hard or coNP-hard unless NP = coNP.*

Proof. Let $\mathcal{R}$ be a relation in TFNP and assume $\mathcal{R}$ is NP-hard (coNP-hard). Then, there is a deterministic polynomial-time Turing machine M that can decide an NP-complete (coNP-complete) language $\mathcal{L}$ when given access to an oracle that on input $I \in \mathcal{I}$ returns a solution $s \in \mathcal{S}$ with $(I, s) \in \mathcal{R}$. We can build a non-deterministic Turing machine M' that decides $\overline{\mathcal{L}}$ as follows: M' simulates M until it accesses the oracle with some query $I \in \mathcal{I}$. It guesses the answer $s \in \mathcal{S}$ of the oracle and it can verify its guess in polynomial-time since $\mathcal{R}$ is in FNP. If the guess was wrong, M' rejects. Otherwise, M' continues to simulate M in this fashion until M would either accept or reject and then inverts the answer from M. Because P is in TFNP, there always exists some $s \in \mathcal{S}$ with $(I, s) \in \mathcal{R}$ and thus there always exists a correct guess for the answer of the oracle. Then, M' accepts if and only if M rejects and hence M' decides $\overline{\mathcal{L}}$. ■

In particular, unless NP = coNP, Theorem 2.1 implies that all of the above presented decision problems are strictly harder than computing any solution to the corresponding search problems.

The class TFNP is called a *semantic class* since the promise that a relation in TFNP is total is not a consequence of the syntactic problem definition. There are no known complete problems for TFNP and this is attributed to the empiric observation that in general "semantic classes seem to have no complete problems" [61, page 499]. In [61], Papadimitriou proposes to subdivide TFNP into syntactic classes that each capture a certain existence proof technique. At this point (even before TFNP was introduced), the class *polynomial-time local search* (PLS) [39] was already defined. PLS captures computational problems corresponding to existence theorems that can be proven with potential arguments. The pigeonhole principle is captured by the class PPP, and several parity arguments by the classes PPAD, PPA, PPADS [61], and a combination of both potential and parity argument is captured in CLS [28]. Here, only one related problem was considered so far: the problem

GIVEN d pairs of points $C_1, \ldots, C_d \subset \mathbb{Q}^d$ and a colorful choice that embraces the origin,

FIND another colorful choice that embraces the origin,

was shown to be PPAD-complete by Meunier and Sarrabezolles [55, Proposition 2], which was posed as an open problem in [53]. However, again this does not bound the computational complexity of COLORFULCARATHÉODORY.

2.1. Overview

The application of Sarkaria's lemma in the reduction from TVERBERG and COLORFULKIRCHBERGER to COLORFULCARATHÉODORY creates highly degenerate instances. In Chapter 3, we discuss how to construct equivalent COLORFULCARATHÉODORY instances that satisfy several general position assumptions. In Chapter 4, we show that the cone version of COLORFULCARATHÉODORY is in the complexity class PPAD. The proof is based on a new topological proof of the colorful Carathéodory theorem that uses Sperner's Lemma [24] by Meunier and Sarrabezolles [70]. Using linear programming techniques, we show how to replace nonconstructive parts of this proof by algorithms. In Chapter 5, we give a formulation of the cone version of COLORFULCARATHÉODORY as a PLS-problem. Finally, we show that a slight change in the definition of COLORFULCARATHÉODORY results in a PLS-complete problem, the *nearest colorful polytope problem*.

Please note that all algorithms in this part of the thesis are analyzed in the WORD-RAM model with logarithmic costs.

Remark 2.2. Bárány and Onn [11] showed that a special case of the convex version of COLORFULCARATHÉODORY can be solved in polynomial time: if each of the $d+1$ color classes contains a ball of radius $\rho > 0$ that is centered at the origin in its convex hull, and if

- all points in $\bigcup_{i=1}^{d+1} C_i$ have ℓ_2-norm between 1 and 2, and
- $1/\rho$ is polynomial in the length of the input,

then a $\mathbf{0}$-embracing colorful choice can be computed in polynomial time. Please note that we consider the cone version of COLORFULCARATHÉODORY instead of the convex version and more importantly, we have no general position assumptions. ◀

2.2. Preliminaries

2.2.1. FNP Reductions

Since FNP contains relations and not languages, a different concept of reduction is necessary to define complete problems. We say an FNP problem A is *FNP-reducible* to an FNP problem B if there exist two polynomial-time computable functions $f_{A \mapsto B}$ and $f_{B \mapsto A}$ with the following properties. Let $\mathcal{I}_A$ denote the set of instances of A and let $\mathcal{S}_A$ denote the set of candidate solutions of A. Define $\mathcal{I}_B$ and $\mathcal{S}_B$ similarly. The function $f_{A\mapsto B}: \mathcal{I}_A \mapsto \mathcal{I}_B$ maps problem instances of A to problem instances of B. The second function $f_{B\mapsto A}: \mathcal{I}_A \times (\mathcal{S}_B \cup \{\perp\}) \mapsto (\mathcal{S}_A \cup \{\perp\})$ maps candidate solutions of B to candidate solutions of A such that the following holds: let $s_B \in \mathcal{S}_B \cup \{\perp\}$ be a candidate solution of B with $\left(f_{A\mapsto B}(I_A), s_B\right) \in B$ or $\perp$ if no such candidate solution exists. If I_A is solvable, then $\left(I_A, f_{B\mapsto A}(I_A, s_B)\right) \in A$ and otherwise $f_{B\mapsto A}(I_A, s_B) = \perp$. The existence of these two functions implies that any polynomial-time algorithm for B can be extended to a polynomial-time algorithm for A.

2.2.2. The Complexity Class PPAD

The complexity class *polynomial parity argument in a directed graph* (PPAD) [61] 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. See Figure 2.1 for an example.

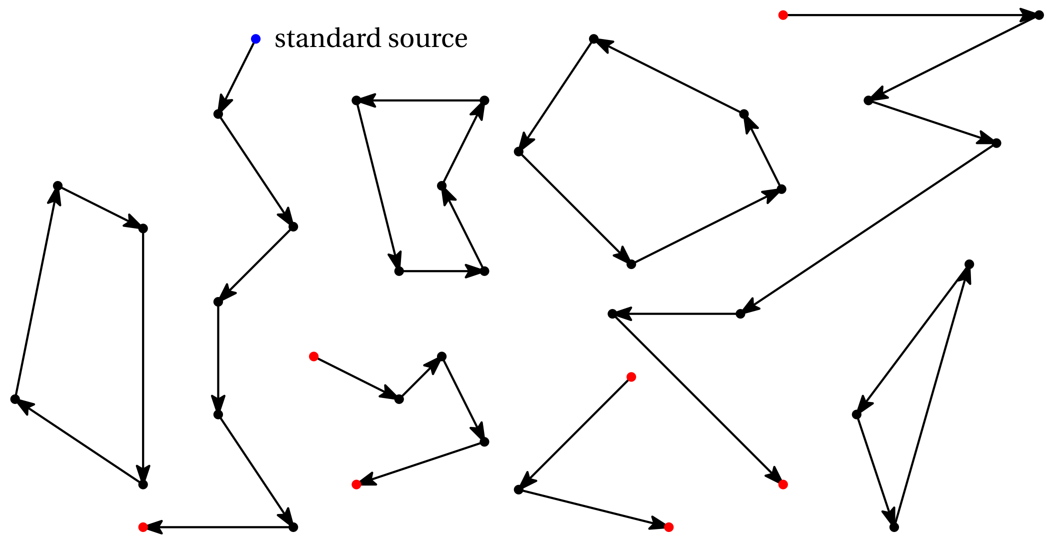

Figure 2.1.: An example of a PPAD-graph problem. The given source is colored blue and we call it the *standard source*. The red nodes are the solutions.

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 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 $\text{pred}(s) = \perp$. 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 $v = \text{succ}(u)$ and $u = \text{pred}(v)$. 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.

To define PPAD-hard and -complete problems, the same concept of reductions as for FNP is used. A problem A is PPAD-hard if all problems in PPAD can be FNP-reduced to A and A is PPAD-complete if A is PPAD-hard and $A \in$ PPAD. The list of PPAD-complete problems includes, among others, computational problems corresponding to several fixed point theorems such as Sperner's lemma in two dimensions [21], the Borsuk-Ulam theorem [61], Brouwer's fixed point theorem [61], as well as computing mixed Nash equilibria in 2-player games [20] and approximating mixed Nash equilibria in k-player games for $k \geq 3$ [27, Theorem 12], [19].

2.2.3. The Complexity Class PLS

The complexity class *polynomial-time local search* (PLS) [1,39,56] 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 $\text{poly}(|I|)$ 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. See Figure 2.2 for an example.

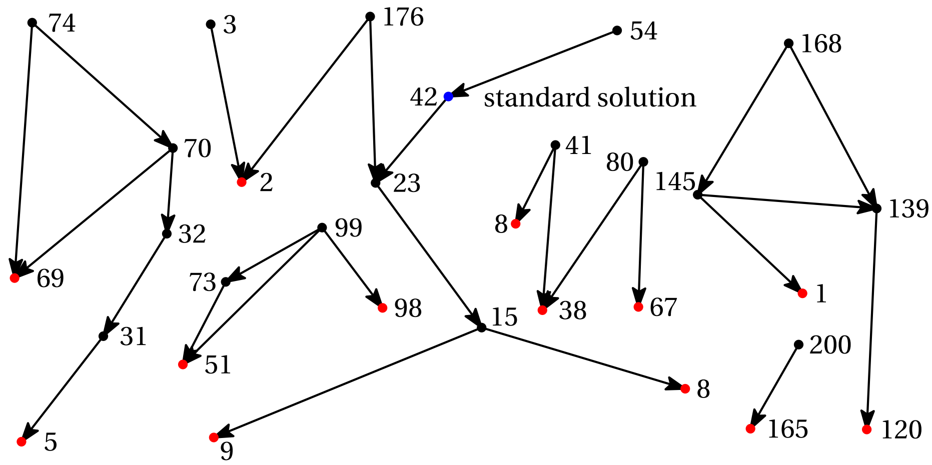

Figure 2.2.: An example of a PLS-graph for a minimization problem. The blue node is the standard solution, and the red nodes denote local optima. The numbers next to the nodes denote the costs.

We say a problem $A \in$ FNP is PLS-hard if all problems in PLS can be FNP-reduced to A and we say A is PLS-complete if $A \in$ PLS and A is PLS-hard. The canonical PLS-complete problem is FLIP [39, Theorem 1]: given a Boolean circuit of polynomial size with n inputs and m outputs, find an input-assignment such that the resulting output interpreted as a number in binary cannot be decreased by flipping one bit in the input. The set of PLS-complete problems includes, among various local search variants and heuristics for NP-complete problems, the Lin-Kernighan heuristic for the traveling salesman problem [60], computing stable configurations in Hopfield neuronal networks [72, Corollary 5.12], and computing pure Nash equilibria in congestion games [32, Theorem 3].

2.2.4. Linear Programming

In this section, we briefly repeat the terminology of linear programming and establish a common notation.

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

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

where $\boldsymbol{b} \in \mathbb{Q}^d$ and $A \in \mathbb{Q}^{d\times n}$ with $\mathrm{rank}(A) = k$. By multiplying with the least common denominator, we may assume in the following that A and $\boldsymbol{b}$ have only integer entries. 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, the coordinates of the corresponding basic solution $\boldsymbol{x}$ are not 0 in dimensions $\mathrm{ind}(B)$. 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 that we obtain by extending the linear system L' with the constraints $\boldsymbol{x} \geq \boldsymbol{0}$ and with a cost vector $\boldsymbol{c} \in \mathbb{Q}^n$:

$$\begin{aligned} L: \quad & \min \boldsymbol{c}^T\boldsymbol{x} \\ \text{s.t. } & A\boldsymbol{x} = \boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0}. \end{aligned}$$

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 basic solution, i.e., let $\boldsymbol{x}$ be such that $A\boldsymbol{x} = \boldsymbol{b}$ and $(\boldsymbol{x})_i = 0$ for $i \in [n] \setminus \mathrm{ind}(B)$. We call $\boldsymbol{x}$ a *basic feasible solution* if $\boldsymbol{x} \geq \boldsymbol{0}$, and we call B a *feasible basis* if the corresponding basic solution is feasible. Furthermore, we say L is *non-degenerate* if for all feasible bases B, the coordinates of the corresponding basic feasible solution are strictly greater than 0. In the following, we assume L is non-degenerate. Now, let $R = [n] \setminus \mathrm{ind}(B)$ denote the set of indices from columns in A that are 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

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

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 [23]. 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 $\mathrm{ind}(B)$ 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

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

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 $\mathbb{R}^n$. Let $H_n^- = \left\{h_{i\geq 0} \subset \mathbb{R}^n \mid i \in [n]\right\}$ denote the set of n coordinate halfspaces, where $h_{i\geq 0} = \{\boldsymbol{x} \in \mathbb{R}^n \mid (\boldsymbol{x})_i \geq 0\}$ and let $H_{L'}$ denote the set of d hyperplanes that are defined by the rows of $A\boldsymbol{x} = \boldsymbol{b}$. Then, all feasible solutions for L are contained in the polyhedron

$$\mathcal{P} = \left(\bigcap_{h^- \in H_n^-} h^-\right) \cap \left(\bigcap_{h \in H_{L'}} h\right).$$

Since L is non-degenerate, $\mathcal{P}$ is simple. Let $f \subseteq \mathcal{P}$ be a k-face of $\mathcal{P}$. Then, f is contained in the intersection of the d hyperplanes from $H_{L'}$ with $n-k$ boundary hyperplanes $(\boldsymbol{x})_{j_1} = 0$, $(\boldsymbol{x})_{j_2} = 0, \ldots, (\boldsymbol{x})_{j_{n-k}} = 0$ from H_n, where $j_1, \ldots, j_{n-k} \in [n]$. The feasible solutions in f can only vary in the k dimensions with indices $[n] \setminus \{j_1, \ldots, j_{n-k}\}$. This set of k dimensions is unique for the face f and we use it to encode f combinatorially:

$$\operatorname{supp}\left(f\right) = [n] \setminus \left\{j_1, \ldots, j_{n-k}\right\}.$$

We call $\operatorname{supp}\left(f\right)$ the *support* of f and we say the columns in $A_{\operatorname{supp}(f)}$ *define* f. In particular, f projected onto dimensions $\operatorname{supp}\left(f\right)$ is the polyhedron that is defined by the linear program $A_{\operatorname{supp}(f)}\boldsymbol{x}' = \boldsymbol{b}$, $\boldsymbol{x}' \geq \boldsymbol{0}$. Furthermore, note that for all subfaces $\check{f} \subseteq f$, we have $\operatorname{supp}\left(\check{f}\right) \subseteq \operatorname{supp}\left(f\right)$ and in particular, all bases that define vertices of f are d-subsets of columns from $A_{\operatorname{supp}(f)}$.

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

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

We conclude this section with the following standard lemma that bounds the number of bits that is necessary to represent a basic feasible solutions for a linear program (e.g., see [73, Corollary 3.2d] for a similar statement).

Lemma 2.3. *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 [50], we can express the ith coordinate of $\boldsymbol{x}_{\mathrm{ind}(B)}$ 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:

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

And similarly, $\left|\det A'_i\right| \leq N$ 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. ■

3 Equivalent Instances of ColorfulCarathéodory 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 ColorfulCarathéodory instances by extending known perturbation techniques for linear programming to our setting. More formally, let $I = (C_1, \ldots, C_d, \boldsymbol{b})$ 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 $C_1^{\approx}, \ldots, C_d^{\approx} \subset \mathbb{Z}^d$ 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) ***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 [26, Section 10-2] and [51] but extends to a more general setting in which the matrix is also perturbed. Based on these results, we then show in Section 3.2 how to construct ColorfulCarathéodory instances with properties (P1)–(P3).

3.1. Polynomials with Bounded Integer Coefficients

In the following, we consider equation systems

$$L_\varepsilon : A\boldsymbol{x} = \boldsymbol{b}, \tag{3.1}$$

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 3.1. *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

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

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 3.1 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 3.2. *Let $L_\varepsilon : A\boldsymbol{x} = \boldsymbol{b}$ be a system of equations as defined in (3.1) 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

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

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, $c_{i,j} \in \mathbb{P}\left[M'\right]$. 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 $k = \max\left\{\deg b_i \mid i \in [d]\right\}$. Then, Lemma 3.1 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. ■

3.2. Construction

Let $C'_1, \dots, C'_d \subset \mathbb{Q}^d$ be d sets that ray-embrace $\boldsymbol{b}' \in \mathbb{Q}^d$. By applying Theorem 1.1, we can ensure that $|C'_i| \le d$ 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 $\|\boldsymbol{b}\|_1 \ge \|\boldsymbol{p}\|_1$ 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 3.2 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 3.3. *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 $(C_1(\varepsilon) \dots C_d(\varepsilon))$. 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 3.2 implies that L_ε is non-degenerate for all $\varepsilon \in (0, \frac{1}{2M})$, 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 3.3 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 3.4. *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 $\boldsymbol{p} + \frac{1}{s}\boldsymbol{m}_{\varepsilon_0} \in \operatorname{pos}\left(P_{\varepsilon_0}(\boldsymbol{p})\right)$ 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

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

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 $\boldsymbol{p} + \frac{1}{s}\boldsymbol{m}_{\varepsilon_0} \in \operatorname{conv}\left(P_{\varepsilon_0}(\boldsymbol{p})\right) \subset \operatorname{pos}\left(P_{\varepsilon_0}(\boldsymbol{p})\right)$, as claimed. ■

As a consequence of Lemma 3.3, 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 3.5. *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 $C' = \left\{\boldsymbol{p} \,\middle|\, i \in [d], \boldsymbol{c}_i \in P_{\varepsilon_0}(\boldsymbol{p})\right\}$ 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 3.3, 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 $\mathbf{0}$ 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 3.6. *Let* $I = \left(C'_1, \ldots, C'_d, \boldsymbol{b}'\right)$ *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* $I^{\approx} = \left(C^{\approx}_1, \ldots, C^{\approx}_d, \boldsymbol{b}^{\approx}\right)$ *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 3.4, each color class $C_i(\varepsilon_0)$ ray-embraces $\boldsymbol{b}(\varepsilon_0)$, so we can apply Theorem 1.1 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 $C^{\approx}_i$ 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 3.3, 3.4, and 3.5. ■

4 ColorfulCarathéodory is in PPAD

We begin by sketching a nonconstructive proof of the colorful Carathéodory theorem that uses Sperner's Lemma. In Section 4.1, we present the main steps of the formulation as a PPAD-problem. A section is dedicated to each major step, and we finally combine these results in Section 4.4 to the main theorem of this chapter.

We briefly restate the definition of a polyhedral complex before we proceed. We call a finite set of polyhedra $\mathcal{P}$ in $\mathbb{R}^d$ a *polyhedral complex* if and only if

- for all polyhedra $f \in \mathcal{P}$, all faces of f are contained in $\mathcal{P}$, and
- for all $f, f' \in \mathcal{P}$, the intersection $f \cap f'$ is a face of both.

Note that the first property 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, please refer to [82, Section 5.1].

The following proof was already published in the PhD thesis of Sarrabezolles [70] and we sketch here the main steps to provide some intuition on the approach that we are taking when we cast ColorfulCarathéodory as PPAD-problem.

Let $\Delta^{d-1} = \operatorname{conv}(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d) \subset \mathbb{R}^d$ denote the standard $(d-1)$-simplex, where $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d$ is the canonical basis of $\mathbb{R}^d$. Furthermore, let $\mathcal{S}$ be a simplicial subdivision of Δ^{d-1} and let $\lambda : V(\mathcal{S}) \mapsto [d]$ be a function that assigns to each vertex $\boldsymbol{v}$ of a simplex in $\mathcal{S}$ a label from $[d]$ such that if $\boldsymbol{v} \in \operatorname{conv}(\boldsymbol{e}_{i_1}, \ldots, \boldsymbol{e}_{i_k})$, then $\lambda(\boldsymbol{v}) \in \{i_1, \ldots, i_k\}$. We call a labeling with this property a *Sperner labeling*. For a simplex $\sigma = \operatorname{conv}\{\boldsymbol{v}_0, \ldots, \boldsymbol{v}_{k-1}\} \in \mathcal{S}$, we denote with $\lambda(\sigma) = \{\lambda(\boldsymbol{v}_i) \mid i \in [k-1]_0\}$ the set of labels of its vertices. Furthermore, we call σ *almost fully-labeled* if $[k-1] \subseteq \lambda(\sigma)$ and we call it *fully-labeled* if $[k] = \lambda(\sigma)$.

Lemma 4.1 (Sperner's Lemma (Strong Version) [24])**.** *The number of fully-labeled $(d-1)$-simplices is odd.*

Proof. For $i \in [d]$, we denote with Σ_i the set

$$\Sigma_i = \left\{\sigma \in \mathcal{S} \,\middle|\, \sigma \subseteq \operatorname{conv}(\boldsymbol{e}_1, \ldots, \boldsymbol{e}_i),\ \sigma \text{ is an almost fully-labeled } (i-1)\text{-simplex}\right\} \subseteq \mathcal{S}.$$

Note that $\Sigma_1 = \{\boldsymbol{e}_1\}$. We define a graph $G = (V, E)$ with $V = \bigcup_{i=1}^d \Sigma_i$ as follows. There is an undirected edge between to nodes $\sigma, \sigma' \in V$ if and only if (i) $\sigma, \sigma' \in \Sigma_i$, $i \in [d]$, and both simplices share a facet $\breve{\sigma}$ with $\lambda(\breve{\sigma}) = [i-1]$, or (ii) one simplex, say σ, is contained in Σ_i for some $i \in [d-1]$, σ' is contained in Σ_{i+1}, and σ is a facet of σ'. See Figure 4.1 for an example in two dimensions.

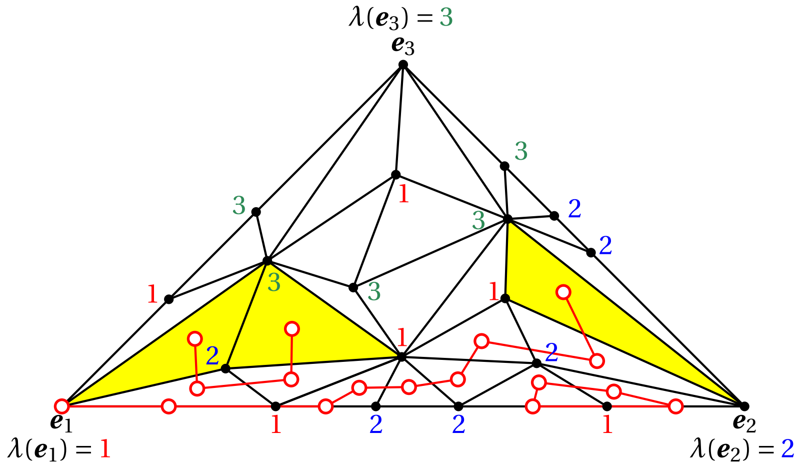

Figure 4.1.: An example of Sperner's lemma in two dimensions. Fully-labeled 2-simplices are colored yellow, and the graph that is defined in the proof of Lemma 4.1 is marked red.

Let $\Sigma_d^\star \subseteq \Sigma_d$ be the set of fully-labeled simplices and set $D_1 = \Sigma_1 \cup \Sigma_d^\star$. We claim that the degree of a node $\sigma \in V$ is 1 if and only if $\sigma \in D_1$ and otherwise 2. First we observe that $\{\boldsymbol{e}_1\}$ is the single node in Σ_1 and hence $\deg\{\boldsymbol{e}_1\} \leq 1$. Second, since $\mathcal{S}$ is a simplicial subdivision of Δ^{d-1}, there exists a 1-simplex $\sigma \subseteq \operatorname{conv}(\boldsymbol{e}_1, \boldsymbol{e}_2)$ with $\{\boldsymbol{e}_1\}$ as its facet. Then, $\sigma \in \Sigma_1$ and hence $\deg\{\boldsymbol{e}_1\} = 1$. Let now $\sigma \in \Sigma_d^\star$ be a fully-labeled $(d-1)$-simplex. Then, σ has exactly one facet $\breve{\sigma}$ with $\lambda(\breve{\sigma}) = [d-1]$ that is either shared with another simplex in Σ_d or is itself a simplex in Σ_{d-1}. Thus, $\deg \sigma = 1$.

Consider now a simplex $\sigma \in \Sigma_i \setminus D_1$, $i \in [d]$. If $\lambda(\sigma) = [i]$, then there exists exactly one simplex in Σ_{i+1} with σ as a facet. Furthermore, σ has exactly one facet with label set $[i-1]$ and this facet is either shared with another simplex in Σ_i or it is itself a simplex in Σ_{i-1}. Hence, $\deg \sigma = 2$. Otherwise, if $\lambda(\sigma) = [i-1]$, σ has exactly two facets with label set $[i-1]$ and each facet is either shared with another simplex in Σ_i or is itself a simplex in Σ_{i-1}.

Because there is an even number of nodes with odd degree, $|D_1|$ must be even. Furthermore, since $\{\boldsymbol{e}_1\}$ is the only simplex in D_1 that is not fully-labeled, the number of fully-labeled simplices in $\mathcal{S}$ is odd. ■

In particular, there is at least one fully-labeled simplex.

We now prove the cone version of Theorem 1.2 using Lemma 4.1. Let $C_1, \ldots, C_d \subset \mathbb{R}^d$ be d color classes of d vectors each that ray-embrace a point $\boldsymbol{b} \in \mathbb{R}^d$, $\boldsymbol{b} \neq \boldsymbol{0}$. Then, let $A = (C_1\ C_2\ \ldots\ C_d) \in \mathbb{R}^{d \times d^2}$ 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. We denote with L^{CC} the linear system

$$L^{\mathrm{CC}}: A\boldsymbol{x} = \boldsymbol{b} \qquad \boldsymbol{x} \geq \boldsymbol{0}. \tag{4.1}$$

Let $\mathcal{P}^{\mathrm{CC}} \subset \mathbb{R}^{d^2}$ denote the polyhedron that is defined by L^{CC}. We note that each color class C_i, $i \in [d]$, is a feasible basis of L^{CC} and hence defines a vertex $\boldsymbol{v}_i$ of $\mathcal{P}^{\mathrm{CC}}$. We call a feasible basis B of L^{CC} *colorful* if B is a colorful choice of $C_1, \ldots, C_d$. Similarly, we call the vertex $\boldsymbol{v}_B$ that is defined by a feasible basis B *colorful* if B is colorful. Now, we cut out a polyhedral complex $\mathcal{F}$ of $\mathcal{P}^{\mathrm{CC}}$ that behaves similarly to a simplex and we show that for an arbitrary triangulation of $\mathcal{F}$ and an appropriate Sperner labeling, all fully-labeled simplices have a colorful vertex. Then, Lemma 4.1 guarantees the existence of a colorful basis.

For a subset $I \subseteq [d]$ of the colors, we denote with L_I^{CC} the following linear subsystem of L^{CC} that is restricted to the colors in I:

$$\begin{aligned} L_I^{\mathrm{CC}}:\ A_{\mathrm{ind}(C_I)}\boldsymbol{x} &= \boldsymbol{b} \\ \boldsymbol{x} &\geq \boldsymbol{0}, \end{aligned}$$

where $C_I = \bigcup_{i \in I} C_i$. As described in Section 2.2.4, the polyhedron that is defined by L_I^{CC} is the projection of a face $f_I \subseteq \mathcal{P}^{\mathrm{CC}}$. Now, we define for each k-subset of colors $I \subseteq [d]$ a $k-1$-dimensional polyhedral subcomplex $\mathcal{F}_I$ of f_I by induction on k. We begin with $k = 1$ and set $\mathcal{F}_{\{i\}} = \{\emptyset, f_{\{i\}}\} = \{\emptyset, \boldsymbol{v}_i\}$ for $i \in [d]$. Let now $k \geq 2$ and let $I \subseteq [d]$ be a set of size k. Assume we already have defined $\mathcal{F}_{\breve{I}}$ for all $\breve{I} \subset I$ with $|\breve{I}| = k-1$. Then, we set $\breve{\mathcal{F}}_I = \bigcup_{\breve{I} \subset I : |\breve{I}| = k-1} \mathcal{F}_{\breve{I}}$ to be the collection of all $(k-2)$-dimensional polyhedral complexes that are defined by subsets of I. Note that $\breve{\mathcal{F}}_I \subset f_I$ since $f_{\breve{I}} \subset f_I$ for all $(k-1)$-subsets $\breve{I}$ of I. By connectedness, there exists a $(k-1)$-dimensional polyhedral subcomplex $\mathcal{F}_I$ of f_I with $\breve{\mathcal{F}}_I$ as boundary. We denote with $\mathcal{F}$ the $(d-1)$-dimensional polyhedral subcomplex $\mathcal{F}_{[d]}$ of $\mathcal{P}^{\mathrm{CC}}$. To apply Lemma 4.1, we fix an arbitrary triangulation $\mathcal{T}$ of $\mathcal{F}$ and we define for each vertex $\boldsymbol{v}$ of $\mathcal{F}$ a label $\lambda(\boldsymbol{v}) \in [d]$ as follows:

$$\lambda(\boldsymbol{v}) = \operatorname*{arg\,max}_{i \in [d]} \left| \operatorname{ind}(C_i) \cap \operatorname{supp}(\boldsymbol{v}) \right|. \tag{4.2}$$

In case of a tie, we take the smallest $i \in [d]$ that achieves the maximum. Now, for a k-set $I \subseteq [d]$, we identify the $(k-1)$-face $g = \operatorname{conv}\{\boldsymbol{e}_i \mid i \in I\}$ of the standard $(d-1)$-simplex Δ^{d-1} with the $(k-1)$-face $f_I \in \mathcal{F}$. Since all bases that define vertices in f_I are subsets of $\bigcup_{i \in I} C_i$, the label of a vertex in f_I is from I. Even though the polyhedral complex $\mathcal{F}$ may not necessarily be homeomorphic to Δ^{d-1}, it still behaves almost like a simplex and the proof of Lemma 4.1 can be mimicked on $\mathcal{F}$ to show that there exists at least one fully-labeled $(d-1)$-simplex in $\mathcal{T}$. We conclude the proof sketch of Theorem 1.2 by showing that a fully-labeled simplex implies a colorful feasible basis.

Lemma 4.2. *Let $\sigma \in \mathcal{T}$ be a fully-labeled simplex. Then, one of the bases that define the vertices of σ is colorful.*

Proof. By construction, $\mathcal{F}$ is a $(d-1)$-dimensional polytopal subcomplex of $\mathcal{P}^{\mathrm{CC}}$. Hence, σ lies on a $(d-1)$-face f of $\mathcal{P}^{\mathrm{CC}}$. The face f is defined by $d + (d-1) = 2d-1$ columns from A. Since there are d colors, one color, say $i^\star$, appears at most once in f. As the simplex σ is fully-labeled, it has a vertex with label $i^\star$ and let $B^\star$ be the feasible basis that defines this vertex. By definition of the labeling λ, the basis $B^\star$ has then at most one vector from each color class, so $B^\star$ is a colorful feasible basis. Note that this in particular implies $i^\star = 1$. ■

4.1. Outline

Inspired by the presented proof of the colorful Carathéodory theorem using Sperner's lemma, we cast COLORFULCARATHÉODORY as a PPAD-problem by replacing nonconstructive steps with algorithms. As before, we would like to navigate in a polyhedral subcomplex $\mathcal{F}$ of $\mathcal{P}^{\mathrm{CC}}$ that behaves like a simplex and then adapt the constructive proof of Sperner's lemma to find a colorful vertex. However, instead of working directly with $\mathcal{F}$, we define a polytopal subdivision $\mathcal{Q}_\Delta$ of the standard simplex Δ^{d-1} that has a dual relationship to $\mathcal{F}$. The polytopal complex $\mathcal{Q}_\Delta$ has no counterpart in the proof from the previous section. We use it as a proxy for $\mathcal{F}$ that enables us on the one hand to work in an easier geometric setting, namely the standard simplex, while on the other hand, due to the dual relationship to $\mathcal{F}$, it is still possible to attack computational problems that involve $\mathcal{Q}_\Delta$ with linear programming techniques by translating them back to $\mathcal{F}$.

In Section 4.2, we define $\mathcal{F}$ implicitly as a solution space to a family of parametrized linear programs. The linear programs all have the linear system L^{CC} as constraints and differ only in their cost vectors. The cost vectors are defined by a linear function that maps points $\boldsymbol{\mu} \in \mathbb{R}^d$ to a cost vector $\boldsymbol{c}_{\boldsymbol{\mu}} \in \mathbb{R}^{d^2}$. We call $\mathbb{R}^d$ the *parameter space* and $\boldsymbol{\mu} \in \mathbb{R}^d$ a *parameter vector*. To each face f of $\mathcal{P}^{\mathrm{CC}}$, 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 $\mathcal{M} = \left\{\boldsymbol{\mu} \in \mathbb{R}^d \,\middle|\, \boldsymbol{\mu} \geq \mathbf{0}, \|\boldsymbol{\mu}\|_\infty = 1\right\}$ 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 is crucial to define later on a Sperner labeling based on the colors in the support. Now, the polyhedral complex $\mathcal{F}$ consists of all faces f with $\Phi(f) \cap \mathcal{M} \neq \emptyset$. Furthermore, the intersections of the parameter regions with $\mathcal{M}$ induce a polytopal complex $\mathcal{Q}$. By performing a central projection with the origin as center of $\mathcal{Q}$ onto the standard simplex Δ^{d-1}, we obtain the polytopal subdivision $\mathcal{Q}_\Delta$ of Δ^{d-1}.

In Section 4.3, use the barycentric subdivision $\operatorname{sd}\mathcal{Q}_\Delta$ to obtain a simplicial subdivision of the standard simplex. Similar to the labeling of the proof in the previous section, we assign to a vertex $\boldsymbol{v}$ in $\operatorname{sd}\mathcal{Q}_\Delta$ the label $\lambda(\boldsymbol{v}) = i$ if the ith color appears most often in the support of the face in $\mathcal{F}$ that corresponds to this vertex, where ties are broken by taking the smallest label. The color controlling property of our cost function then implies that λ is a Sperner labeling, and a similar argument as in the proof of Lemma 4.2 shows that one of the vertices of a fully-labeled $(d-1)$-simplex corresponds to a colorful feasible basis. This concludes a new constructive variant of the topological proof. 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 $\mathcal{Q}_\Delta$ such that the encoding of simplices that share a facet differs only in one part. Finally, in Section 4.4, we combine the results from the previous sections together with a generalization of the formulation of 2D-Sperner as a PPAD-problem [61] to show that COLORFULCARATHÉODORY is in PPAD.

4.2. The Polytopal Complex

Let $I = (C_1, \ldots, C_d, \boldsymbol{b})$ be a fixed instance of COLORFULCARATHÉODORY. Applying Lemma 3.6, we can assume without loss of generality that I has properties (P1)–(P3).

We begin by defining a family of linear programs $\{L_{\boldsymbol{\mu}}^{\mathrm{CC}} \mid \boldsymbol{\mu} \in \mathbb{R}^d\}$, where each linear program $L_{\boldsymbol{\mu}}^{\mathrm{CC}}$ consists of the linear system L^{CC} (see (4.1)) with respect to the instance I 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 $\boldsymbol{\mu} \in \mathbb{R}^d$. More formally, we denote with $L_{\boldsymbol{\mu}}^{\mathrm{CC}}$ the linear program

$$L_{\boldsymbol{\mu}}^{\mathrm{CC}}: \quad \begin{aligned} \min\ & \boldsymbol{c}_{\boldsymbol{\mu}}^T \boldsymbol{x} \\ \text{s.t. }\ & A\boldsymbol{x} = \boldsymbol{b} \\ & \boldsymbol{x} \geq \boldsymbol{0}, \end{aligned} \tag{4.3}$$

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. Furthermore, we want the cost vector to control the colors that appear in the support of optimal faces when restricted to parameter vectors in $\mathcal{M}$. More precisely, we want that if $(\boldsymbol{\mu})_i = 0$, then *no* optimal basis for $L_{\boldsymbol{\mu}}$ contains a column with color i, while $(\boldsymbol{\mu})_i > 0$ implies that color i *may* appear in an optimal basis. Let $N = d!m^d$ be as defined in Lemma 2.3, where m is the largest absolute value that appears in A and $\boldsymbol{b}$. Then, we define $\boldsymbol{c}_{\boldsymbol{\mu}}$ as

$$(\boldsymbol{c}_{\boldsymbol{\mu}})_j = 1 + \big(1 - (\boldsymbol{\mu})_i\big)\, dN^2 + \varepsilon^j, \tag{4.4}$$

where $j \in \left[d^2\right]$, i is the color of the jth column in A, and $0 < \varepsilon \leq N^{-3}$ is a perturbation that we define shortly. We prove the desired color-controlling properties of the cost function at the end of this section.

In the following, we denote for a face $f \subseteq \mathcal{P}^{\mathrm{CC}}$, $f \neq \emptyset$, with

$$\Phi(f) = \left\{ \boldsymbol{\mu} \in \mathbb{R}^d \,\middle|\, f \text{ is optimal for } L_{\boldsymbol{\mu}} \right\}$$

the set of all parameter vectors for which f is optimal. We call this the *parameter region* for f. By our discussion from Section 2.2.4, 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:

$$\begin{aligned} L_{B,f}^{\Phi}: (\hat{\boldsymbol{r}}_{B,\boldsymbol{c}_{\boldsymbol{\mu}}})_j &= 0 \text{ for } j \in \operatorname{supp}(f) \setminus \operatorname{ind}(B) \\ (\hat{\boldsymbol{r}}_{B,\boldsymbol{c}_{\boldsymbol{\mu}}})_j &\leq 0 \text{ for } \left[d^2\right] \setminus \operatorname{supp}(f). \end{aligned}$$

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

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

It is easy to see that $\mathcal{F} \cup \{\emptyset\}$ is a polyhedral subcomplex of $\mathcal{P}^{\mathrm{CC}}$: first, all elements from $\mathcal{F}$

are faces of $\mathcal{P}^{\mathrm{CC}}$. Second, since $f \in \mathcal{F}$, there exists a parameter vector $\boldsymbol{\mu} \in \mathcal{M}$ for which f is optimal. Then, all nonempty subfaces of f are also optimal for $\boldsymbol{\mu}$ and hence they are contained in $\mathcal{F}$. The intersections of the parameter regions with faces of $\mathcal{M}$ induce a subdivision $\mathcal{Q}$ of $\mathcal{M}$:

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

In the remaining part of this section, we show that $\mathcal{Q}$ is a $(d-1)$-dimensional polytopal complex and characterize its elements. We begin with a perturbation lemma to ensure that the linear system $L^{\Phi}_{B,f}$ extended by the constraints $\boldsymbol{\mu} \in g$, where g is a face of $\mathcal{M}$, is non-degenerate for all faces f of $\mathcal{P}^{\mathrm{CC}}$ and all choices of bases B for a vertex of f.

Lemma 4.3. *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 $H_\Phi = \left\{h_j \mid j \in \left[d^2\right] \setminus \operatorname{ind}(B)\right\}$. Furthermore, let $H_\square$ denote the set of supporting hyperplanes for the facets of the unit cube in $\mathbb{R}^d$. Then, for all k-subsets H' of $H_\Phi \cup H_\square$, 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 $H_\Phi \cup H_\square$, 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 $H'_\square = H' \cap H_\square$ the hyperplanes from $H_\square$. Set $R = \left[d^2\right] \setminus \operatorname{ind}(B)$ 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

$$\left(\left(\boldsymbol{c}_{\boldsymbol{\mu}}\right)_R - \left(A^{-1}_{\operatorname{ind}(B)} A_R\right)^T \left(\boldsymbol{c}_{\boldsymbol{\mu}}\right)_{\operatorname{ind}(B)}\right)_{\operatorname{rank}_R(\phi_i)} = 0 \text{ for } i \in [n], \tag{4.5}$$

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, (4.5) is equivalent to

$$-dN^2(\boldsymbol{\mu})_{\operatorname{col}(\phi_i)} + dN^2\boldsymbol{a}_i^T \begin{pmatrix} (\boldsymbol{\mu})_{\operatorname{col}(\beta_1)} \\ \vdots \\ (\boldsymbol{\mu})_{\operatorname{col}(\beta_d)} \end{pmatrix} = -1 - dN^2 - \varepsilon^{\phi_i} + \boldsymbol{a}_i^T \begin{pmatrix} 1 + dN^2 + \varepsilon^{\beta_1} \\ \vdots \\ 1 + dN^2 + \varepsilon^{\beta_d} \end{pmatrix} \text{ for } i \in [n], \tag{4.6}$$

where $\operatorname{col}\left(\phi_i\right)$ and $\operatorname{col}\left(\beta_i\right)$ denote the colors of the columns with indices ϕ_i and β_i, respectively. Thus, (4.6) is of the form

$$A_\Phi \boldsymbol{\mu} = \boldsymbol{b}_\Phi, \tag{4.7}$$

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 2.3, the fractions in $A^{-1}_{\operatorname{ind}(B)}$ all have the same denominator: $\det A_{\operatorname{ind}(B)} \in \mathbb{Z}$. We set $A'_\Phi = \left(\det A_{\mathrm{ind}(B)}\right) A_\Phi$ and $\boldsymbol{b}'_\Phi = \left(\det A_{\mathrm{ind}(B)}\right) \boldsymbol{b}_\Phi$. Then, the linear system

$$A'_\Phi \boldsymbol{\mu} = \boldsymbol{b}'_\Phi \tag{4.8}$$

is equivalent to (4.7), 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 2.3 the absolute value of the entries in $A^{-1}_{\mathrm{ind}(B)}$ is at most N, there exists a constant $c' \in \mathbb{N}$ such that $m' \le N^{c'}$ and c' is independent of the choice of B.

Set $n' = \left|H'_\square\right|$. Since we assume that the hyperplanes in H' have a point in common and since $H'_\square \subseteq H'$, the hyperplanes in $H'_\square$ 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 (4.8), we can express the intersection of hyperplanes in H' as

$$\left(A'_\Phi\right)_{[d]\setminus J} (\boldsymbol{\mu})_{[d]\setminus J} = \boldsymbol{b}'_\Phi - \sum_{j\in J_1} (A'_\Phi)_j. \tag{4.9}$$

The matrix $\left(A'_\Phi\right)_J$ 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 3.2 implies that for all $\varepsilon \in \left(0, \frac{1}{2M}\right)$, the right hand vector of (4.9) 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 (4.9) has a solution, it follows that the rank of (4.9) 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 $\mathcal{Q}$ are indeed polytopes and by characterizing precisely their dimension and their facets.

Lemma 4.4. *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*

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

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

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

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 4.3, 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 4.3, we have $H_g^= \cap H_{\Phi(f)}^= = \emptyset$, the set $H^=$ contains $d - \dim g + \dim f$ hyperplanes. Again by Lemma 4.3, 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 $h_1^-, \ldots, h_{j^\star-1}^-$, 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 4.3.

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 $\breve{g} = g \cap h$ is a facet of g and we have $\check{q} = \Phi(f) \cap \breve{g}$. Assume now $h^- \in H_{\Phi(f)}^-$ and let h be defined by the equation $(\hat{\boldsymbol{r}}_{B,\boldsymbol{c}_\mu})_j = 0$ 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 $\breve{g}$ be a facet of g with $\check{q} = \Phi(f) \cap \breve{g} \neq \emptyset$. 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

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

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 4.4 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 4.4 immediately bounds the maximum dimensions of faces in $\mathcal{F}$ by $d-1$.

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

Lemma 4.5. *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,*

$$q_1 \cap q_2 = \Phi\big(\hat{f}\big) \cap \breve{g},$$

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

Proof. We begin with showing that $\Phi(f_1) \cap \Phi(f_2) = \Phi\big(\hat{f}\big)$. 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 $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$ and thus $\Phi(f_1) \cap \Phi(f_2) \subseteq \Phi\big(\hat{f}\big)$. Let now $\boldsymbol{\mu}$ be a parameter vector from $\Phi\big(\hat{f}\big)$. 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, $\Phi\big(\hat{f}\big) = \Phi(f_1) \cap \Phi(f_2)$. Then, we can express $q_1 \cap q_2$ as

$$q_1 \cap q_2 = \big(\Phi(f_1) \cap g_1\big) \cap \big(\Phi(f_2) \cap g_2\big) = \Phi\big(\hat{f}\big) \cap \breve{g},$$

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

Equipped with Lemmas 4.4 and 4.5, we are now ready to show that $\mathcal{Q}$ is a polytopal complex.

Lemma 4.6. *The set $\mathcal{Q}$ is a $(d-1)$-dimensional polytopal complex that decomposes $\mathcal{M}$.*

Proof. Lemma 4.4 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 4.4, 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 4.5 to express the intersection of q_1 and q_2 as $\Phi\big(\hat{f}\big) \cap \breve{g}$. Since $\hat{f} \in \mathcal{F}$ and since $\breve{g}$ 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 $\breve{g}$ is a face of g_1, a repeated application of Lemma 4.4 shows that $q_1 \cap q_2$ is a face of q_1. Similarly, because $\hat{f}$ is a superface of f_2 and

$\breve{g}$ is a face of g_2, a repeated application of Lemma 4.4 proves that $q_1 \cap q_2$ is a face of q_2, as desired. ∎

A further implication of Lemmas 4.4 and 4.5 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 4.7. *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 4.5, we can write q as $\Phi\left(\hat{f}\right) \cap \breve{g}$, where $\hat{f} \in \mathcal{F}$ is the smallest face in $\mathcal{P}^{\mathrm{CC}}$ that contains f_1 and f_2 and $\breve{g}$ is a face of g_1 and of g_2. If $\hat{f} \neq f_1$ or $\breve{g} \neq g_1$, then by Lemma 4.4,

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

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

From now on, we denote the $(d-1)$-dimensional standard simplex with $\Delta \subset \mathbb{R}^d$, and we perform a central projection with the origin as center of $\mathcal{Q}$ onto Δ. It is easy to see that this projection is a bijection. For a parameter vector $\boldsymbol{\mu} \in \mathcal{M}$, we denote with

$$\Delta(\boldsymbol{\mu}) = \frac{1}{\|\boldsymbol{\mu}\|_1}\boldsymbol{\mu}$$

its projection onto Δ. Similarly, for a parameter vector $\boldsymbol{\mu} \in \Delta$, we denote with

$$\mathcal{M}(\boldsymbol{\mu}) = \frac{1}{\|\boldsymbol{\mu}\|_\infty}\boldsymbol{\mu}$$

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 4.1 for an overview of the current and future notation. The following results are immediate consequences from Lemmas 4.4, 4.6, 4.7, and the fact that the projection is a bijection.

Corollary 4.8 (of Lemma 4.4)**.** *Let $q \neq \emptyset$ be an element from $\mathcal{Q}_\Delta$. Then, there exists a face $f \in \mathcal{F}$ and a face $g \in \mathcal{S}$ 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*

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

∎

Corollary 4.9 (of Lemma 4.6)**.** *The set $\mathcal{Q}_\Delta$ is a $(d-1)$-dimensional polytopal complex that decomposes Δ.* ∎

Corollary 4.10 (of Lemma 4.7)**.** *Let $q \in \mathcal{Q}_\Delta$ be a polytope. Then, there exists unique pair of faces f, g with $f \in \mathcal{F}$ and $g \in \mathcal{S}$ such that $q = \Phi_\Delta(f) \cap g$.* ∎

We conclude this section by showing the claimed color-controlling property of the cost function. Before we proceed with the proof, we first observe the following direct implication of Property (P2).

Observation 4.11. *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.*

Lemma 4.12. *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 $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$. Then, $B^\star \cap C_{i^\times} = \emptyset$.*

Proof. Let $\boldsymbol{x}^\star$ be the basic feasible solution for $B^\star$ with respect to $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$. 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 4.11 and Lemma 2.3, we have $\left(\boldsymbol{x}^\star\right)_k \geq 1/N$. Hence,

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

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 2.3, $\left(\boldsymbol{x}^{(i^\star)}\right)_j$ 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 $B^\star$. ∎

4.3. The Barycentric Subdivision

The *barycentric subdivision* [49, Definition 1.7.2] is a well-known method to subdivide a polytopal complex into simplices. The barycentric subdivision $\operatorname{sd}\mathcal{Q}_\Delta$ of $\mathcal{Q}_\Delta$ consists of all simplices $\operatorname{conv}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_{k-1})$ such that there exists a chain $q_0 \subset \cdots \subset q_{k-1}$ of polytopes in $\mathcal{Q}_\Delta$ with $\dim q_{i-1} < \dim q_i$ and $\boldsymbol{v}_i$ being the barycenter of q_i for $i \in [k-1]_0$ and $k \in [d]$. Similarly to the labeling from (4.2), we define the label of a vertex $\boldsymbol{v} \in \operatorname{sd}\mathcal{Q}_\Delta$ as follows. By Corollary 4.10, there exists a unique pair $f \in \mathcal{F}$ and $g \in \mathcal{S}$ with $\boldsymbol{v} \in \operatorname{relint}\left(\Phi_\Delta(f) \cap g\right)$. Then, the label $\lambda(\boldsymbol{v})$ of $\boldsymbol{v}$ is defined as

$$\lambda(\boldsymbol{v}) = \underset{i \in [d]}{\arg\max} \left|\operatorname{ind}(C_i) \cap \operatorname{supp}(f)\right|. \tag{4.10}$$

Again, in case of a tie, we take the smallest $i \in [d]$ that achieves the maximum. Lemma 4.12 implies that λ is a Sperner labeling of $\operatorname{sd}\mathcal{Q}_\Delta$: first, consider a vertex $\boldsymbol{e}_i$ of Δ, where $i \in [d]$.

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 (4.4).
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 4.1). $L^{\mathrm{CC}}_{\boldsymbol{\mu}}$ 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 4.12).
$\mathcal{F}$	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 (4.10).
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 (4.11)).

Table 4.1.: Notation reference.

Lemma 4.12 states that the unique optimal feasible basis for $L^{\text{CC}}_{\mathcal{M}(\boldsymbol{e}_i)} = L^{\text{CC}}_{\boldsymbol{e}_i}$ is C_i and thus $\lambda(\boldsymbol{e}_i) = i$. Second, consider a vertex $\boldsymbol{v} \in \operatorname{sd}\mathcal{Q}_\Delta$ and let $\Delta_I = \operatorname{conv}\{\bigcup_{i\in I}\boldsymbol{e}_i\}$ be a face of Δ that contains $\boldsymbol{v}$, where $I \subseteq [d]$. Then, $(\boldsymbol{v})_j = 0$ and hence $(\mathcal{M}(\boldsymbol{v}))_j = 0$ for $j \notin I$. Thus by Lemma 4.12, $\lambda(\boldsymbol{v}) \in I$. Note that λ is in fact a Sperner labeling for any fixed simplicial subdivision of Δ. Now, Lemma 4.1 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 4.13. *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 $\boldsymbol{b}$.*

Proof. Let $q_0 \subset \cdots \subset q_{d-1}$ be the chain that corresponds to σ in $\operatorname{sd}\mathcal{Q}_\Delta$. By Corollary 4.10, 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, Corollary 4.8 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 $\dim q_{d-1} = d-1$, Corollary 4.8 implies that $\dim f_i = d-1-i$ and hence $\left|\operatorname{supp}(f_i)\right| = 2d-1-i$ 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 $\left|\operatorname{supp}(f_{i-1})\right| = \left|\operatorname{supp}(f_i)\right| + 1$ 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 (4.10)). 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 $\left|\operatorname{ind}(C_{i^\star}) \cap \operatorname{supp}(f_{d-1})\right| \geq 2$. Since there is at most one column with color $i^\times$ among $A_{a_0}, \ldots, A_{a_{d-2}}$, we have $\left|\operatorname{supp}(f_i) \cap \operatorname{ind}(C_{i^\times})\right| \leq 1$ 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 $\left|\operatorname{ind}(C_{i^\star}) \cap \operatorname{supp}(f_{d-1})\right| \geq 2$, we have $\lambda(f_i) \neq i^\times$ for all $i \in [d-1]_0$, a contradiction to σ being fully-labeled. ■

This concludes a new constructive proof of the colorful Carathéodory theorem. However, in order to show that COLORFULCARATHÉODORY is contained in PPAD, we need to replace the invocation of Lemma 4.1 to find a fully-labeled $(d-1)$-simplex in $\mathcal{Q}_\Delta$ by a PPAD-problem. Note that it is not possible to use the formulation of Sperner from [61, 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 [61] to $\mathcal{Q}_\Delta$ by mimicking the proof of Lemma 4.1. For this, we need to be able to find algorithmically simplices in $\operatorname{sd}\mathcal{Q}_\Delta$ that share a given facet. We begin with a simple combinatorial 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 Corollary 4.10, 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. For $j = 0,1$, let $I_j \subseteq [d]$, denote the indices of the coordinates that are fixed to j. Then, the encoding $\mathrm{enc}\left(q\right)$ of q is defined as

$$\mathrm{enc}\left(q\right) = \left(\mathrm{supp}\left(f\right), I_0, I_1\right).$$

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 $\mathrm{enc}\left(\sigma\right)$ is defined as

$$\mathrm{enc}\left(\sigma\right) = \left(\mathrm{enc}\left(q_0\right), \ldots, \mathrm{enc}\left(q_k\right)\right). \tag{4.11}$$

In the proof of Lemma 4.1, 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 Σ_k denote the set

$$\Sigma_k = \left\{\sigma \in \mathrm{sd}\,\mathcal{Q}_\Delta \,\middle|\, \dim(\sigma) = k-1,\ \sigma \subseteq \Delta_{[k]}\right\}$$

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 $g\left(I_0, I_1\right) = \left\{\boldsymbol{\mu} \in \mathcal{M} \,\middle|\, j = 0,1,\ (\boldsymbol{\mu})_i = j \text{ for } i \in I_j\right\}$ 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 $Q_i = \left(S^{(i)}, I_0^{(i)}, I_1^{(i)}\right)$, $S^{(i)} \subseteq \left[d^2\right]$, and $I_0^{(i)}, I_1^{(i)}$ are disjoint subsets of $[d]$ for $i \in [k-1]_0$. We say T is a *valid k-tuple* 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 $a_{i-1} \in \left[d^2\right] \setminus S^{(i)}$,

  (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 $I_1^{(i-1)} = I_1^{(i)} \cup \left\{j_{i-1}\right\}$, or $I_1^{(i-1)} = I_1^{(i)}$ and $I_0^{(i-1)} = I_0^{(i)} \cup \left\{j_{i-1}\right\}$.

Lemma 4.14. *For $k \in [d]$, the function* enc *restricted to the simplices in Σ_k is a bijection from Σ_k to the set of valid k-tuples.*

Proof. We begin by showing that the encoding $\mathrm{enc}\left(\sigma\right)$ 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 Corollary 4.10, 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, Corollary 4.8 implies that $\dim g_{k-1} = k-1$ and $\dim f_{k-1} = 0$. In particular, $\mathrm{supp}\left(f_{k-1}\right)$ 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, Corollary 4.8 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, (4.11) together with Corollary 4.10 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, Corollary 4.8 states that $\dim q_{k-1} = k-1$. Moreover by Corollary 4.8 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 Corollary 4.8, 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 $\operatorname{enc}(\sigma) = T$. ■

The next lemma shows that simplices that share facets have a similar encoding.

Lemma 4.15. *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 $\operatorname{enc}\left(q_{k-1}\right) = \left(S^{(k-1)}, I_0^{(k-1)}, I_1^{(k-1)}\right)$ and $\operatorname{enc}\left(\hat{q}_k\right) = \left(S^{(k)}, I_0^{(k)}, I_1^{(k)}\right)$. 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 4.14 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

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

as claimed. ∎

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

Proof. 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. ∎

We conclude this section by showing how to traverse Σ efficiently via the respective encodings.

Lemma 4.17. *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:*

- *Given* $\operatorname{enc}(\sigma)$ *and i, compute the encoding of the simplex $\sigma' \in \Sigma_k$ that shares the facet* $\operatorname{conv}\left\{\boldsymbol{v}_j \mid j \in [k-1]_0, j \neq i\right\}$ *with σ or state that there is none;*
- *Assuming that $k < d$ and given* $\operatorname{enc}(\sigma)$*, compute the encoding of the simplex $\hat{\sigma} \in \Sigma_{k+1}$ that has σ as facet;*
- *Assuming that $k > 1$ and given* $\operatorname{enc}(\sigma)$*, compute the encoding of the simplex $\check{\sigma} \in \Sigma_{k-1}$ that is a facet of σ or state that there is none.*

Proof. We begin with the first problem. By Lemma 4.15, 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 4.16. Furthermore, Lemma 4.15 directly implies polynomial-time algorithms for the second and third problem. ■

4.4. The PPAD graph

Using our tools from the previous sections, we now describe the PPAD graph $G = (V, E)$ for the ColorfulCarathéodory instance. The orientation of the edges is postponed until the end of this section. The definition of G follows mainly the ideas from the formulation of Sperner as a PPAD-problem [61, Theorem 2] and the proof of Lemma 4.1.

The graph has one node per simplex in Σ that is almost fully-labeled or fully-labeled. That is, we have one node for each $(k-1)$-simplex in Σ_k whose vertices have as labels $1, \ldots, k$ or $1, \ldots, k-1$. 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 labels but the largest possible label. More formally, let $\sigma = \operatorname{conv}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_{k-1}) \in \operatorname{sd} \mathcal{Q}_\Delta$ be some simplex. Then, we denote with

$$\lambda(\sigma) = \bigcup_{i=0}^{k-1} \{\lambda(\boldsymbol{v}_i)\}$$

the collection of the labels of its vertices and we say σ is $[i]$-labeled if $\lambda(\sigma) = [i]$ for some $i \in [d]$. For $k \in [d]$, we denote with

$$V_k = \{\operatorname{enc}(\sigma) \mid \sigma \in \Sigma_k, [k-1] \subseteq \lambda(\sigma)\}$$

the set of all encodings for $(k-1)$-simplices in Σ_k whose vertices have all labels 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 $\operatorname{enc}(\sigma), \operatorname{enc}(\sigma')$ be two vertices in V_k for some $k \in [d]$. Then, there is an edge between $\operatorname{enc}(\sigma)$ and $\operatorname{enc}(\sigma')$ 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 $\operatorname{enc}(\sigma) \in V_k$ and $\operatorname{enc}(\sigma') \in V_{k+1}$ for some $k \in [d-1]$. Then, there is an edge between $\operatorname{enc}(\sigma)$ and $\operatorname{enc}(\sigma')$ if $\lambda(\sigma) = [k]$ and σ is a facet of σ'. In the next lemma, we show that the underlying undirected graph of G consists only of paths and cycles.

Lemma 4.18. *Let* $\operatorname{enc}(\sigma) \in V$ *be a node. If* $\operatorname{enc}(\sigma) \in V_1$ *or* $\operatorname{enc}(\sigma) \in V_d$ *with* $\lambda(\sigma) = [d]$*, then* $\deg \operatorname{enc}(\sigma) = 1$*. Otherwise,* $\deg \operatorname{enc}(\sigma) = 2$.

Proof. 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 $\deg \mathrm{enc}(\sigma) = 2$ 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 $\deg \mathrm{enc}(\sigma) = 2$ 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 4.14 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 $\mathrm{enc}(\sigma) \in V$, 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 [61].

Let $\mathrm{enc}(\sigma), \mathrm{enc}(\sigma') \in V_d$ be two adjacent nodes. By definition, the encoded simplices $\sigma = \mathrm{conv}(\boldsymbol{v}_0, \dots, \boldsymbol{v}_{d-1})$ and σ' share a facet $\check{\sigma} = \mathrm{conv}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_{d-1})$ 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

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

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 4.18.

Lemma 4.19. *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

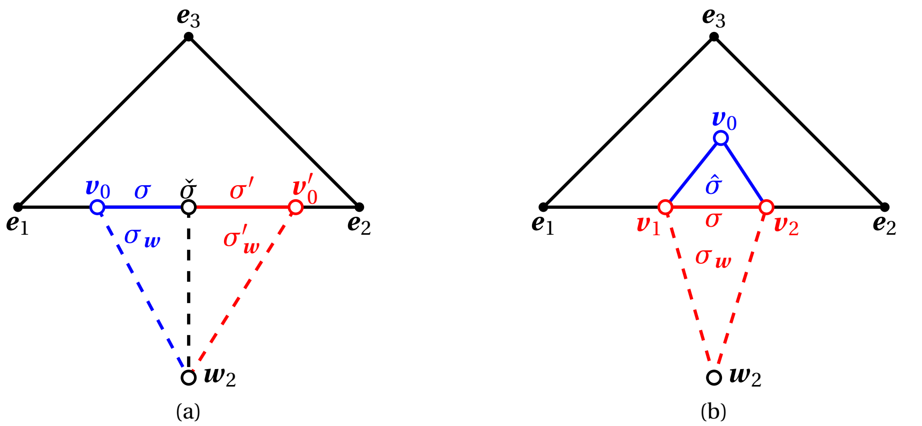

Figure 4.2.: (a) Two simplices of the same dimension are lifted with the same set of vertices $\boldsymbol{w}_i$. (b) The simplices have different dimensions. The lower dimensional simplex is lifted with an additional vertex.

the encoded simplices with $\lambda(\boldsymbol{v}_i) = i$ for $i \in [k-1]$. That is, σ and σ' share the facet $\check{\sigma} = \operatorname{conv}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{k-1})$. 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 $\operatorname{aff}(\Delta_{[k]})$ by the $(k-2)$-dimensional affine space $\operatorname{aff}(\check{\sigma})$. Since $\boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_d \notin \operatorname{aff}(\Delta_{[k]})$, the two vertices $\boldsymbol{v}_0$ and $\boldsymbol{v}'_0$ are separated in $\mathbb{R}^d$ by the hyperplane $\operatorname{aff}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{k-1}, \boldsymbol{w}_{k+1}, \ldots, \boldsymbol{w}_d)$. The situation is depicted in Figure 4.2 (a). Then, we have $\operatorname{dir}(\sigma, \sigma') = -\operatorname{dir}(\sigma', \sigma)$, since

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

Let now $\operatorname{enc}(\sigma) \in V_{k-1}$ and $\operatorname{enc}(\hat{\sigma}) \in V_k$ 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 $\sigma = \operatorname{conv}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{k-1})$ 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 $i \in [k-1]$. See Figure 4.2 (b). 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 $\boldsymbol{w}_{k+1},\dots,\boldsymbol{w}_k \notin \operatorname{aff}(\Delta_{[k]})$, $\boldsymbol{v}_0$ 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

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

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 4.20. *Let $q_0,\dots,q_{k-1} \subset \mathbb{R}^d$ be k polytopes such that $q_0 \subset \dots \subset q_{k-1}$ 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 $\boldsymbol{v}'_0,\dots,\boldsymbol{v}'_{k-1}$ be $k-1$ vectors such that $\boldsymbol{v}'_i \in q_i$ and $\operatorname{aff}(\boldsymbol{v}'_0,\dots,\boldsymbol{v}'_i) = \operatorname{aff}(q_i)$ for all $i \in [k-1]_0$. Then,*

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

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 $\operatorname{aff}(q_i) = \operatorname{aff}(\boldsymbol{v}'_0,\dots,\boldsymbol{v}'_i)$ 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 $\boldsymbol{v}_0,\dots,\boldsymbol{v}_j$ 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 $\operatorname{aff}(\boldsymbol{v}'_0,\dots,\boldsymbol{v}'_{i-1})$. Since both $\boldsymbol{v}_i$ and $\boldsymbol{v}'_i$ lie on the same side of $\operatorname{aff}(\boldsymbol{v}'_0,\dots,\boldsymbol{v}'_{i-1})$ 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, $\boldsymbol{v}'_0,\dots,\boldsymbol{v}'_{i-1} \in \operatorname{aff}(\boldsymbol{v}_0,\dots,\boldsymbol{v}_{i-1})$ 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 4.21. *Let* $\operatorname{enc}(\sigma) = \left(\operatorname{enc}(q_0), \dots, \operatorname{enc}(q_{k-1})\right) \in V$ *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 4.16, 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 $\operatorname{aff}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_i) = \operatorname{aff}(q_i)$. ∎

The following lemma is now an immediate consequence of Lemmas 4.20 and 4.21.

Lemma 4.22. *Let* $\operatorname{enc}(\sigma), \operatorname{enc}(\sigma) \in V$ *be two adjacent nodes. Then, we can compute* $\operatorname{dir}(\sigma, \sigma')$ *in polynomial time.* ∎

With the tools from the last sections, little is left to show that COLORFULCARATHÉODORY is in PPAD.

Theorem 4.23. COLORFULCARATHÉODORY, CENTERPOINT, TVERBERG, SIMPLICIALCENTER, *and* COLORFULKIRCHBERGER *are in PPAD.*

Proof. We give a formulation of COLORFULCARATHÉODORY as PPAD-problem. Lemma 1.14 then 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 $\boldsymbol{b} \neq \boldsymbol{0}$. Let $I^{\approx} = (C_1^{\approx}, \ldots, C_d^{\approx}, \boldsymbol{b}^{\approx})$ denote then the COLORFUL-CARATHÉODORY instance that we obtain by applying Lemma 3.6 to I. 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 $d \in \mathbb{N}$. The set of valid candidate solutions $\mathcal{S}_I$ for the instance I consists of all valid k-tuples with respect to the instance $I^{\approx}$ that encode fully-labeled and almost fully-labeled simplices, as well as all colorful choices with respect to I that ray-embrace $\boldsymbol{b}$, where $k \in [d]$.

The set $\mathcal{S}_I$ is polynomial-time recognizable: let $s \in \mathcal{S}$ be a candidate solution. If it is a tuple, we first use the algorithm from Lemma 4.16 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 4.17 and Lemma 4.22. However, to ensure that the sources and sinks of the graph do not only encode colorful choices with respect to I that ray-embrace $\boldsymbol{b}$ but actually are those colorful choices, we modify the computation of the predecessor and the successor as follows. If a node $s \in V$ is a source different from the standard source in the graph G, it encodes by Lemmas 4.13, 4.18, and 4.19 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. Note that since $I^{\approx}$ has Property (P3) by Lemma 3.6, 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 COLORFULCARATHÉODORY is in PLS

After having established that COLORFULCARATHÉODORY and its descendants are contained in PPAD, we continue with showing that these problems are in fact contained in the intersection of PPAD with the complexity class PLS [39]. In Section 5.2, we further show that already a slight modification of COLORFULCARATHÉODORY results in a PLS-complete problem.

5.1. A PLS Formulation of COLORFULCARATHÉODORY

The proof of the colorful Carathéodory theorem by Bárány as presented in Chapter 1 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 $\boldsymbol{p}^\star \in \text{pos}(C)$ 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

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

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 $O\big(\text{poly}(d,n)L\big)$ time, where L is the length of the quadratic program in binary [40, 44].

Lemma 5.1. *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 $\boldsymbol{p}^\star \in \text{pos}(C)$ with minimum ℓ_2-distance to $\boldsymbol{b}$ in $O\big(\text{poly}(d)L\big)$ time.*

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

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

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

$$\boldsymbol{b}' = \left(\begin{array}{c} 0 \\ \vdots \\ 0 \\ \hline 1 \end{array}\right) \in \mathbb{Q}^{d+1}.$$

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

$$A\boldsymbol{x} = \boldsymbol{b}', \, \boldsymbol{x} \geq \boldsymbol{0} \tag{5.1}$$

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

$$\boldsymbol{x} = \left(\begin{array}{cccccccc|ccc|c} x_1^+ & x_1^- & x_2^+ & x_2^- & \ldots & x_d^+ & x_d^- & & \psi_1 & \ldots & \psi_d & x_{\boldsymbol{b}} \end{array}\right)^T \in \mathbb{R}^{3d+1},$$

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 (5.1), we have $x_{\boldsymbol{b}} = 1$ and thus for $i \in [d]$, the ith equality of (5.1) is equivalent to

$$x_i^+ - x_i^- = (\boldsymbol{p})_i - (\boldsymbol{b})_i. \tag{5.2}$$

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 (5.2), we have

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

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

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

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 $\text{pos}(C)$ with minimum ℓ_2-distance to $\boldsymbol{b}$. Using the algorithm from [40] or [44], we can compute $\boldsymbol{p}^\star$ in $O\big(\text{poly}(d)L\big)$ 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 5.2. *The problems* COLORFULCARATHÉODORY, CENTERPOINT, TVERBERG, SIMPLICIALCENTER, *and* COLORFULKIRCHBERGER *are in PPAD∩ PLS.*

Proof. By Theorem 4.23, COLORFULCARATHÉODORY is in PPAD. We now give a formulation of COLORFULCARATHÉODORY as a PLS-problem. Then statement is then implied by Lemma 1.14.

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 $\text{pos}(s)$ to $\boldsymbol{b}$. Using the algorithm from Lemma 5.1, 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. ■

5.2. A PLS-Complete Generalization of COLORFULCARATHÉODORY

Even more interestingly in the light of Theorem 5.2, we can show that a generalization of COLORFULCARATHÉODORY, the *local search nearest colorful polytope problem* (L-NCP) that is motivated by Bárány's original proof [9], results in a PLS-complete problem. Additionally, using a variant of the PLS-completeness proof, we prove that finding a global optimum for NCP (G-NCP) is NP-hard. This answers a question by Bárány and Onn [11, p. 561]. We note that this question has been answered independently by Meunier and Sarrabezolles [55, Theorem 2].

Let $C_1,\ldots,C_m \subset \mathbb{Q}^d$ be m color classes that do not necessarily embrace the origin. For a given set $C' \subset \mathbb{Q}^d$, let $\delta(C') = \min\left\{\|\boldsymbol{c}\|_1 \mid \boldsymbol{c} \in \operatorname{conv}(C')\right\}$ denote the minimum ℓ_1-norm of a point in $\operatorname{conv}(C')$. In L-NCP, we want to find a colorful choice C such that $\delta(C)$ cannot be decreased by swapping a single point with another point of the same color. In the global search variant G-NCP, we want to find a colorful choice C such that $\delta(C)$ is minimum.

In the language of PLS, L-NCP is defined as follows.

Definition 5.3 (L-NCP)**.**

Instances. The set of problem instances $\mathcal{I}$ consists of all tuples $(C_1,\ldots,C_m)$, where $d \in \mathbb{N}$ and for $i \in [m]$, we have $C_i \subset \mathbb{Q}^d$.

Candidate solutions. The set of candidate solutions consists of all sets $C \subset \mathbb{Q}^d$, where $d \in \mathbb{N}$. For a fixed instance $I = (C_1,\ldots,C_m) \in \mathcal{I}$, we define the set of valid candidate solutions $\mathcal{S}_I$ of I to be the set of all colorful choices with respect to $C_1,\ldots,C_m$.

Cost function. Let $s \in \mathcal{S}_I$ be a colorful choice. Then, the cost $c_{I,s}$ of s with respect to I is defined as $\delta(s)$. We want to minimize the costs.

Neighborhood. Let $I \in \mathcal{I}$ be an instance and let $s \in \mathcal{S}_I$ be a valid candidate solution. Then, the set of neighbors $N_{I,s}$ of s consists of all colorful choices that can be obtained by swapping one point with another point of the same color in s.

We reduce the PLS-problem MAX-2SAT/FLIP [72] to L-NCP. In MAX-2SAT/FLIP, we are given a 2-CNF formula, i.e., a Boolean formula in conjunctive normal form in which each clause consists of at most 2 literals, and for each clause a weight. The problem is to find an assignment such that the weighted sum of unsatisfied clauses cannot be decreased by flipping the value of one variable. More formally, MAX-2SAT/FLIP is defined as follows.

Definition 5.4 (MAX-2SAT/FLIP)**.**

Instances. The set of instances $\mathcal{I}'$ consists of all tuples $I = (n, K_1,\ldots,K_m)$ such that $n \in \mathbb{N}$ and for $i \in [n]$, the tuple K_i has the form (w_i, T_i, F_i), where $w_i \in \mathbb{Z}$ and $T_i, F_i \subseteq [n]$ with $|T_i \cup F_i| \le 2$ for all $i \in [n]$. Then, we identify with K_i the clause $\widehat{K}_i = \left(\bigvee_{j\in T_j} x_j\right) \vee \left(\bigvee_{j\in F_j} \overline{x}_j\right)$ with weight w_i and we identify with I the 2-CNF formula $\widehat{K}_1 \wedge \cdots \wedge \widehat{K}_m$ with variables $x_1,\ldots,x_n$.

Candidate solutions. The set of candidate solutions $\mathcal{S}'$ contains all tuples $A = (v_1,\ldots,v_n)$, where $n \in \mathbb{N}$ and $v_i \in \{0,1\}$ for $i \in [n]$. Given an instance $I \in \mathcal{I}'$ in which n variables

$x_1, \ldots, x_n$ appear, we define the set of valid candidate solutions $\mathcal{S}'_I$ for I as the set of all n-tuples from $\mathcal{S}'$. We interpret the ith entry of a tuple $A \in \mathcal{S}'_I$ as an assignment to x_i and we denote it with $A(x_i)$.

Cost function. Let $I \in \mathcal{I}'$ be an instance. Then, we define the cost $c'_{I,s}$ of a valid candidate solution $s \in \mathcal{S}'_I$ as the sum of the weights of all unsatisfied clauses. We want to minimize the cost.

Neighborhood. Let $I \in \mathcal{I}'$ be an instance and $s \in \mathcal{S}'_I$ a tuple of size n. Then, the set of neighbors $N'_{I,s}$ of s consists of all tuples that can be obtained by replacing the $i$the entry $A(x_i)$ with $1 - A(x_i)$, where $i \in [n]$.

The following theorem is due to Schäffer and Yannakakis.

Theorem 5.5 ([72, Corollary 5.12]). MAX-2SAT/FLIP *is PLS-complete.*

We continue with the reduction from MAX-2SAT/FLIP to L-NCP.

Theorem 5.6. L-NCP *is PLS-complete.*

Proof. Let $I' = (n, K_1, \ldots, K_d) \in \mathcal{I}'$ be a fixed instance of MAX-2SAT/FLIP. We construct an instance $I \in \mathcal{I}$ of L-NCP in which each colorful choice C encodes an assignment A_C such that the cost $c_{I,C}$ of C equals the cost c'_{I',A_C}.

For each variable x_i, we introduce a color class $X_i = \{\boldsymbol{x}_i, \overline{\boldsymbol{x}_i}\}$ consisting of two points in $\mathbb{Q}^d$ that encode whether x_i is set to 1 or 0. We assign the jth dimension to the jth clause and set

$$(\boldsymbol{x}_i)_j = \begin{cases} -nw_j & \text{if } x_i = 1 \text{ satisfies } \widehat{K}_j, \text{ and} \\ w_j & \text{otherwise,} \end{cases}$$

where $j \in [d]$. Similarly, we set

$$\left(\overline{\boldsymbol{x}_i}\right)_j = \begin{cases} -nw_j & \text{if } x_i = 0 \text{ satisfies } \widehat{K}_j, \text{ and} \\ w_j & \text{otherwise,} \end{cases}$$

where $j \in [d]$. Then, a colorful choice C of $X_1, \ldots, X_m$ corresponds to the assignment $A_C \in \mathcal{S}'_{I'}$ that sets x_i to 1 if $\boldsymbol{x}_i \in C$ and otherwise to 0.

In the following, we construct an instance of L-NCP such that the convex hull of a colorful choice C contains the origin if projected onto the dimensions corresponding to clauses that are satisfied by A_C (and hence do not contribute to the cost of C). Moreover, if projected onto the subspace corresponding to the unsatisfied clauses, $\delta(C)$ equals the total weight of unsatisfied clauses which then defines completely the cost of C.

We introduce additional helper color classes to decrease the distance to the origin in dimensions that correspond to satisfied clauses. In particular, we have for each clause $\widehat{K}_i$, $i \in [m]$, a color class $H_i = \{\boldsymbol{h}_i\}$ consisting of a single point, where

$$(\boldsymbol{h}_i)_j = \begin{cases} (d+1)\left((n+2) - \frac{d}{d+1}\right) w_i & \text{if } j = i, \text{ and} \\ w_j & \text{otherwise,} \end{cases}$$

where $j \in [d]$. The last helper color class $H_{d+1} = \{\boldsymbol{h}_{d+1}\}$ again contains a single point, but now all coordinates are set to the clause weights, i.e.,

$$(\boldsymbol{h}_{d+1})_j = w_j \text{ for } j \in [d].$$

See Figure 5.1 for an example.

Let now $I = (X_1, \ldots, X_n, H_1, \ldots, H_{d+1}) \in \mathcal{I}$ denote the constructed L-NCP instance. We continue with showing that the cost of a colorful choice equals the cost of the corresponding assignment by proving the following two inequalities.

(i) for every colorful choice $C \in \mathcal{S}_I$, the cost are lower bounded by the cost of the corresponding assignment:
$$c_{I,C} \geq c'_{I',A_C}.$$

(ii) for every colorful choice $C \in \mathcal{S}_I$, the cost are upper bounded by the cost of the corresponding assignment:
$$c_{I,C} \leq c'_{I',A_C}.$$

Note that (i) and (ii) directly imply that L-NCP is PLS-complete. To see this, consider a local optimum $s^\star \in \mathcal{S}_I$ of the L-NCP instance I. By definition, the costs of all other colorful choices that can be obtained from $s^\star$ by swapping one point with another of the same color are greater or equal to $c_{I,s^\star}$. Then, the total weight of unsatisfied clauses by the corresponding assignment $A_{s^\star} \in \mathcal{S}'_{I'}$ cannot be decreased by flipping a variable. Thus, $A_{s^\star}$ is a local minimum of the MAX-2SAT/FLIP instance I'.

(i) Let $C \in \mathcal{S}_I$ be a colorful choice and assume some clause $\widehat{K}_j$ is not satisfied by the corresponding assignment $A_C \in \mathcal{S}'_{I'}$. By construction, the jth coordinate of each point $\boldsymbol{p}$ in C is at least w_j. Thus, the jth coordinate of every convex combination of the points in C is at least w_j and hence $c_{I,C} \geq c_{I',A_C}$.

(ii) Let $C \in \mathcal{S}_I$ be a colorful choice. In the following, we construct a convex combination of the points in C that results in a point $\boldsymbol{p}$ whose ℓ_1-norm is exactly the total weight of unsatisfied clauses in the corresponding assignment $A_C \in \mathcal{S}'_{I'}$ and thus $c_{I,C} \leq c_{I',A_C}$. For $k = 0, 1, 2$, let S_k denote the set of clauses that are satisfied by exactly k literals with respect to the assignment A_C. As a first step towards constructing $\boldsymbol{p}$, we show the existence of an intermediate point in the convex hull of the helper classes.

Lemma 5.7. *There is a point $\boldsymbol{h} \in \operatorname{conv}(H_1, \ldots, H_{d+1})$ whose jth coordinate is $(n+2)w_j$ if $j \in S_2$ and w_j otherwise.*

Proof. Take $\boldsymbol{h} = \sum_{i \in S_2} \frac{1}{d+1} \boldsymbol{h}_i + \left(1 - \frac{|S_2|}{d+1}\right) \boldsymbol{h}_{d+1}$. Then, for $j \in S_0 \cup S_1$, we have

$$(\boldsymbol{h})_j = \sum_{i \in S_2} \frac{1}{d+1} (h_i)_j + \left(1 - \frac{|S_2|}{d+1}\right) (\boldsymbol{h}_{d+1})_j \overset{j \notin S_2}{=} \sum_{i \in S_2} \frac{1}{d+1} w_j + \left(1 - \frac{|S_2|}{d+1}\right) w_j = w_j.$$

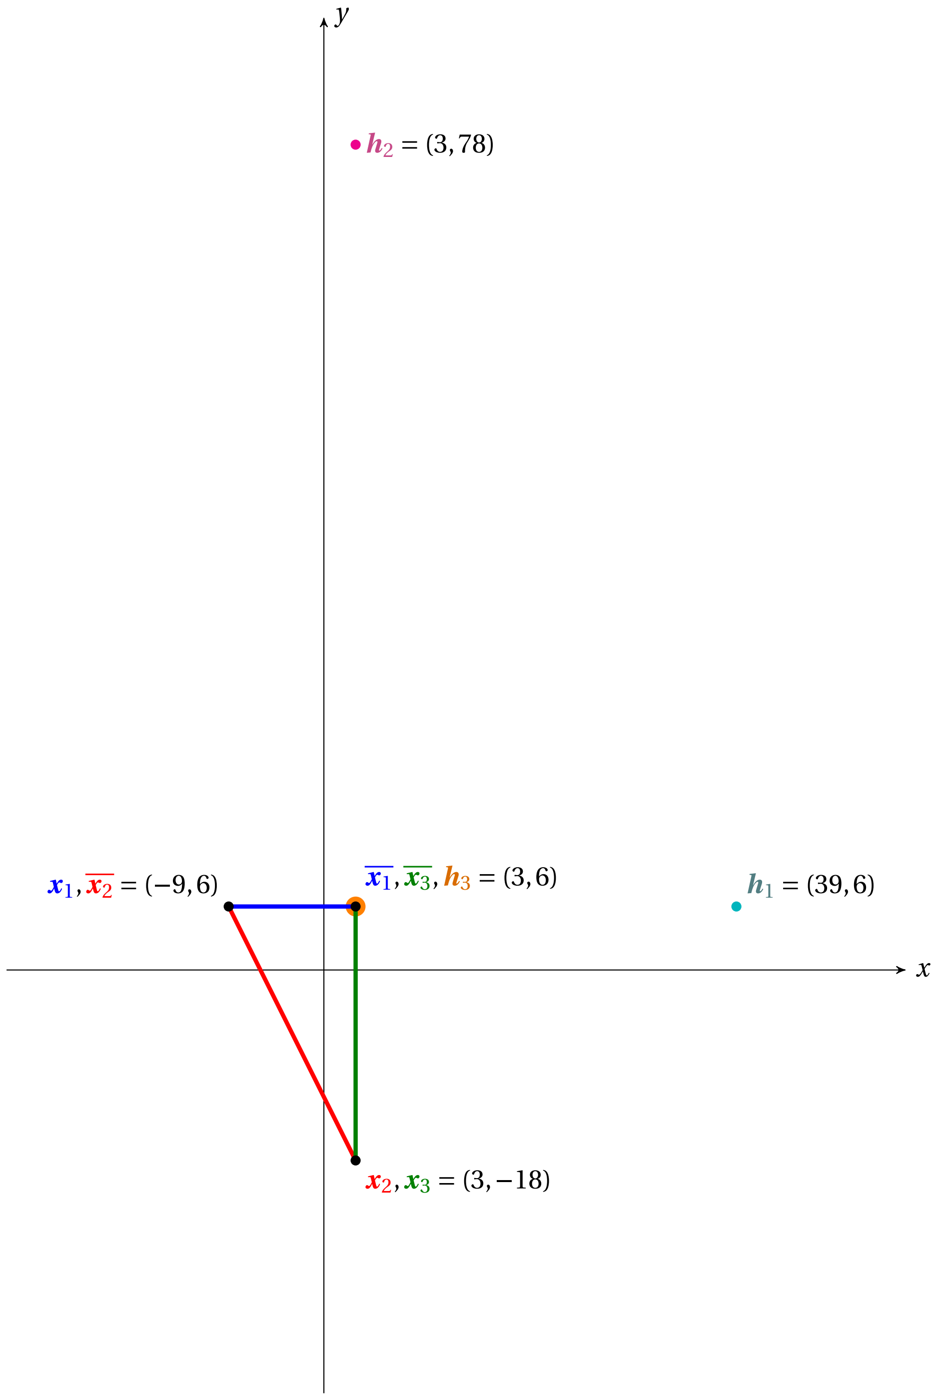

Figure 5.1.: Construction of the point sets corresponding to the Max-2SAT/Flip instance $(x_1 \vee \overline{x_2}) \wedge (x_2 \vee x_3)$ with weights 3 and 6, respectively.

And for $j \in S_2$, we have

$$\begin{aligned}(\boldsymbol{h})_j &= \sum_{i \in S_2} \frac{1}{d+1}(h_i)_j + \left(1 - \frac{|S_2|}{d+1}\right)(\boldsymbol{h}_{d+1})_j \\ &= \frac{1}{d+1}\boldsymbol{h}_j + \sum_{i \in S_2 \setminus \{j\}} \frac{1}{d+1}(h_i)_j + \left(1 - \frac{|S_2|}{d+1}\right)(h_{d+1})_j \\ &= \left((n+2) - \frac{d}{d+1}\right) w_j + \frac{d}{d+1} w_j = (n+2) w_j,\end{aligned}$$

as desired. ■

Let now $\boldsymbol{l}_i$ be the point from X_i in the colorful choice C and consider the point

$$\boldsymbol{p} = \frac{1}{n+1}\left(\sum_{i=1}^{n} \boldsymbol{l}_i + \boldsymbol{h}\right),$$

where $\boldsymbol{h}$ is the point from Lemma 5.7. We show that $(\boldsymbol{p})_j = w_j$ if $j \in S_0$, and otherwise $(\boldsymbol{p})_j = 0$. Let j be an clause index from S_0. Since A_C does not satisfy $\widehat{K}_j$, the jth coordinate of the points $\boldsymbol{l}_1, \dots, \boldsymbol{l}_n$ is w_j. Also, $(\boldsymbol{h})_j = w_j$ by Lemma 5.7. Thus, $(\boldsymbol{p})_j = w_j$. Consider now some clause index $j \in S_1$ and let $b \in [2]$ be the index of the point $\boldsymbol{l}_b$ that corresponds to the single literal that satisfies $\widehat{K}_j$. Then, we have

$$\begin{aligned}(\boldsymbol{p})_j &= \sum_{i=1}^{n} \frac{1}{n+1}(\boldsymbol{l}_i)_j + \frac{1}{n+1}(\boldsymbol{h})_j \\ &= \frac{1}{n+1}(\boldsymbol{l}_b)_j + \sum_{i=1, i \neq b}^{n} \frac{1}{n+1}(\boldsymbol{l}_i)_j + \frac{1}{n+1}(\boldsymbol{h})_j = \frac{-n}{n+1} w_j + \frac{n}{n+1} w_j = 0.\end{aligned}$$

Finally, consider some clause index $j \in S_2$ and let b_1, b_2 be the indices of the two literals that satisfy $\widehat{K}_j$. Then, we obtain

$$\begin{aligned}(\boldsymbol{p})_j &= \sum_{i=1}^{n} \frac{1}{n+1}(\boldsymbol{l}_i)_j + \frac{1}{n+1}(\boldsymbol{h})_j \\ &= \frac{1}{n+1}(\boldsymbol{l}_{b_1})_j + \frac{1}{n+1}(\boldsymbol{l}_{b_2})_j + \sum_{i=1, i \notin \{b_1, b_2\}}^{n} \frac{1}{n+1}(\boldsymbol{l}_i)_j + \frac{1}{n+1}(\boldsymbol{h})_j \\ &= \frac{-2n}{n+1} w_j + \frac{n-2}{n+1} w_j + \frac{n+2}{n+1} w_j = 0,\end{aligned}$$

and thus $\|\boldsymbol{p}\|_1 = c_{I', A_C}$, as claimed. ■

A straightforward modification of the reduction from Theorem 5.6 shows that finding a globally optimal solution for an L-NCP instance is NP-hard by a reduction from 3SAT.

Theorem 5.8. G-NCP *is NP-hard.*

Proof. The proof of Theorem 5.6 can be adapted easily to reduce 3SAT to G-NCP. Given a set of clauses $K_1, \dots, K_d$, we set the weight of each clause to 1 and construct the same point sets as in the PLS reduction. Additionally, we introduce for each clause K_j a new helper color class $H'_j = \{\boldsymbol{h}'_j\}$, where

$$(\boldsymbol{h}'_i)_j = \begin{cases} (d+1)\left((2n+2) - \frac{d}{d+1}\right) & \text{if } i = j, \text{ and} \\ 1 & \text{otherwise.} \end{cases}$$

Let now C be a colorful choice and let A_C be the corresponding assignment. As in the PLS-reduction, for $k = 0, \ldots, 3$, let S_k contain all clauses that are satisfied by exactly k literals in the assignment A_C. Then, the following point $\boldsymbol{h}$ is contained in the convex hull of the helper points:

$$\boldsymbol{h} = \sum_{i \in S_2} \frac{\boldsymbol{h}_i}{d+1} + \sum_{j \in S_3} \frac{\boldsymbol{h}'_j}{d+1} + \left(1 - \frac{|S_2|}{d+1}\right) \boldsymbol{h}_{d+1}.$$

Again, the convex combination $\boldsymbol{p} = \sum_{i=1}^{n} \frac{1}{n+1} \boldsymbol{l}_i + \frac{1}{n+1} \boldsymbol{h}$ results in a point in the convex hull of C whose distance to the origin is the number of unsatisfied clauses, where $\boldsymbol{l}_i$ denotes the point from X_i in C. Together with (i) from the proof of Theorem 5.6, 3SAT can be decided by knowing a global optimum $C^\star$ to the NCP problem: if $\delta(C^\star) = 0$, $A_{C^\star}$ is a satisfying assignment. If not, there exists no satisfying assignment at all. ■

As mentioned in the introduction, we can adapt the proof of Theorem 5.8 to answer a question by Bárány and Onn [11]. Again, this result was obtained independently by Meunier and Sarrabezolles [55, Theorem 2].

Corollary 5.9. *Let $C_1, \ldots, C_m \subset \mathbb{Q}^d$ be an input for* G-NCP*. Then,* G-NCP *remains NP-hard even if $m = d + 1$.*

Proof. Let F be a 3SAT formula with d clauses and n variables. As in the proof of Theorem 5.8, we construct $n + 2d + 1 =: d' + 1$ point sets in $\mathbb{Q}^d$ such that there is a colorful choice that embraces the origin if and only if F is satisfiable. Since $d' > d$, we can lift the point sets to $\mathbb{Q}^{d'}$ by appending 0-coordinates. Then, we have $d' + 1$ point sets such that there is a colorful choice that embraces the origin if and only if F is satisfiable. ■

II

Algorithms for ColorfulCarathéodory

6 Introduction to Part II

Since there are no polynomial-time algorithms for ColorfulCarathéodory known, approximation algorithms are of interest. This was first considered by Bárány and Onn [11] who described how to find a colorful choice whose convex hull is "close" to the origin under several general position assumptions. We call a set ε*-close* to the origin if its convex hull has ℓ_2-distance at most ε to $\mathbf{0}$. Let in the following $\varepsilon, \rho > 0$ be parameters. Given $d+1$ sets $C_1, \dots, C_{d+1} \in \mathbb{Q}^d$ such that

(i) each C_i, $i \in [d+1]$, contains a ball of radius ρ centered at the origin in its convex hull,

(ii) all points $\boldsymbol{p} \in C_i$, $i \in [d+1]$, fulfill $1 \leq \|\boldsymbol{p}\| \leq 2$, and

(iii) the points in all sets can be encoded using L bits.

Then, the algorithm by Bárány and Onn iteratively computes a sequence of colorful choices whose ℓ_2-distances of their convex hulls to the origin strictly decrease until a colorful choice that embraces the origin is found. In particular, if stopped earlier, a colorful choice that is ε-close to $\mathbf{0}$ can be computed in time $\text{poly}(L, \log(1/\varepsilon), 1/\rho)$ on the Word-Ram with logarithmic costs. Note that if $1/\rho = O\big(\text{poly}(L)\big)$, the algorithm actually finds a colorful choice with the origin in its convex hull in polynomial-time. The Bárány-Onn algorithm is essentially the algorithm from the proof of the convex version of Theorem 1.2 and the main contribution is a careful analysis.

In the same spirit, Barman [12] showed that if the points have constant norm, a colorful choice that is ε-close to the origin can be found in $O\left(d^{O(1/\varepsilon^2)} L\right)$ time, where L is the length of the binary encoding of the color classes. The algorithm uses the following approximative version of Carathéodory's theorem as main ingredient: let $P \subset \mathbb{R}^d$ be a $\mathbf{0}$-embracing point set. Then, for any $\varepsilon > 0$, there exists a subset $P' \subseteq P$ of size $c_\varepsilon = O\big(m_P/\varepsilon^2\big)$ that is ε-close to $\mathbf{0}$, where $m_P = \max_{\boldsymbol{p} \in P} \|\boldsymbol{p}\|$. This immediately implies a simple brute-force algorithm: let $C_1, \dots, C_{d+1} \subset \mathbb{Q}^d$ be point sets with $\mathbf{0} \in \text{conv}(C_i)$ for $i \in [d+1]$ and assume all points have constant norm. Let further $C \subseteq \bigcup_{i=1}^{d+1} C_i$ be a $\mathbf{0}$-embracing colorful choice whose existence is guaranteed by Theorem 1.2. Then, the approximative version of Carathéodory's theorem asserts that there is a subset $C' \subseteq C$ of size c_ε that is ε-close to the origin. We can now guess C' by trying out all $\binom{d+1}{c_\varepsilon}$ possibilities for the colors in C' and for each color i, we try all $|C_i|$ possibilities of picking a point with color i. For each choice of C', we can check whether it is ε-close to the origin by solving a convex quadratic program. Solving one convex quadratic program needs $O\big(\text{poly}(d)\, L\big)$ time [40, 44]. Hence, assuming that each color class is of size $O(d)$, we can compute ε-close colorful choice in $O\left(d^{O(1/\varepsilon^2)} L\right)$ time.

Both approaches relax the requirement that the computed colorful choice embraces the origin. However, to apply Sarkaria's lemma (Lemma 1.9), it is crucial that lifted points embrace the origin. If the convex hull is only close to the origin but does not contain it, the proof of Sarkaria's lemma breaks. As this is the main tool of the reductions in Section 1.2 to COLORFUL-CARATHÉODORY, it is not immediate how the two above approximation algorithms can be used to obtain approximation algorithms for the descendants of COLORFULCARATHÉODORY. On the other hand, allowing multiple points from each color class may have a natural interpretation in the reduction. This is true for both the proof of Tverberg's theorem (Theorem 1.8) and the proof of the colorful Kirchberger theorem (Theorem 1.10). In TVERBERG, we are given a set $P \subset \mathbb{Q}^d$ of $n = (m-1)(d+1)+1$ points and we want to find a Tverberg partition $\mathcal{T} = \{T_1, \ldots, T_m\}$ of P. Now, in the reduction to COLORFULCARATHÉODORY, each point $\boldsymbol{p}_i \in P$ is mapped to a color class C_i in $\mathbb{Q}^{n-1}$. The color class C_i consists of m points, where the jth point in C_i represents the choice of assigning $\boldsymbol{p}$ to T_j. A set $C_k \subseteq \bigcup_{i=1}^{n} C_i$ with at most k points from each color class C_i, $i \in [n]$, has a clear interpretation: it encodes a family $\widetilde{\mathcal{T}} = \{\widetilde{T}_1, \ldots, \widetilde{T}_m\}$ of subsets of P such that $\bigcap_{j=1}^{m} \operatorname{conv}(T_j) \neq \emptyset$ and each point $\boldsymbol{p}_i \in P$ may appear in up to k sets. By removing sets from $\widetilde{\mathcal{T}}$ until each point from P appears at most once, we can construct out of $\widetilde{\mathcal{T}}$ a Tverberg partition of P of smaller size. In the case of COLORFULKIRCHBERGER, we are given $n = (m-1)(d+1)+1$ Tverberg partitions $\mathcal{T}_1, \ldots, \mathcal{T}_n$ in $\mathbb{Q}^d$ and we want to merge them into one Tverberg partition by taking one point from each $\mathcal{T}$. Here, each color class C_i in the reduction to COLORFULCARATHÉODORY corresponds to a Tverberg partition $\mathcal{T}_i$ and a set $C_k \subseteq \bigcup_{i=1}^{n} C_i$ with at most k points from each C_i, $i \in [n]$, corresponds to a Tverberg partition $\widetilde{\mathcal{T}}$ that contains at most k points from each $\mathcal{T}_i$. These applications motivate a different approximation problem: given $d+1$ sets $C_1, \ldots, C_{d+1} \subset \mathbb{Q}^d$ that embrace the origin and a parameter k, we want to find a $\mathbf{0}$-embracing set $C \subseteq \bigcup_{i=1}^{d+1} C_i$ such that $|C \cap C_i| \leq k$ for $i \in [n]$. We call such a set a *k-colorful choice* and we want to minimize k. Surprisingly, this notion of approximation has not been studied so far.

6.1. Overview

In Chapter 7, we show that an $\lceil \varepsilon d \rceil$-colorful choice can be computed in polynomial time for any fixed $\varepsilon > 0$. We discuss possible applications of the approximation algorithms to the descendants of COLORFULCARATHÉODORY in Section 7.3.

Moreover, we consider a further notion of approximation that is motivated by our formulation of COLORFULCARATHÉODORY as a PPAD-problem in Chapter 4. Here, we are given only $k \leq d+1$ color classes $C_1, \ldots, C_k \subset \mathbb{Q}^d$ with $\mathbf{0} \in \operatorname{conv}(C_i)$ for $i \in [d+1]$ and we want to find a $\lceil (d+1)/k \rceil$-colorful choice that embraces the origin. Even the case $k = 2$ is already nontrivial and we show how to apply the techniques from Chapter 4 to compute such a set C in weakly polynomial time.

In Chapter 9, we consider the problem of computing an exact solution for COLORFUL-CARATHÉODORY. We show that having a large amount of color classes helps: given color classes $C_1, \ldots, C_n \subset \mathbb{Q}^d$ for $n = \Theta\left(d^2 \log d\right)$ that all embrace the origin, a $\mathbf{0}$-embracing colorful choice can be computed in quasi-polynomial time by a repeated invocation of our tools from Chapter 7.

Finally, we consider in Chapter 10 the problem to compute the simplicial depth of a point with respect to a given point set. We present a new approximation algorithm that computes a $(1+\varepsilon)$-approximation with high probability in time $\widetilde{O}\left(n^{d/2+1}\right)$ if the dimension is fixed, where $\widetilde{O}(\cdot)$ hides polylog-factors. Furthermore, we show that the problem of computing the simplicial depth exactly when the dimension is part of the input becomes $\#\mathsf{P}$-complete [6, 79] and $\mathsf{W}[1]$-hard [25, 31]. This also implies that the problem of computing the colorful simplicial depth is $\mathsf{W}[1]$-hard. Here, only $\#\mathsf{P}$-completeness was known.

Please note that, if not otherwise noted, algorithms in this part of the thesis are analyzed in the REAL-RAM model.

7 k-Colorful Choices

In Section 7.2, we present two approximation algorithms that follow the same strategy. In the first algorithm, we begin with a complete color class and then replace a subset by points from other color classes while maintaining the property that the origin is embraced. The second algorithm performs this replacement step repeatedly, each time decreasing the maximum number of points from a color class, to further improve the approximation guarantee.

The necessary tools to implement the replacement step are presented in the next section. Note that in contrast to the algorithm by Bárány and Onn, we do not assume the input to be in general position but we rather ensure general position by a result similar to Carathéodory's theorem.

7.1. Embracing Equivalent Points

Let $C \subset \mathbb{R}^d$ be a $\mathbf{0}$-embracing point set. We say C is *minimally* $\mathbf{0}$*-embracing* if $C \setminus \{c\}$ is not $\mathbf{0}$-embracing for all points $\boldsymbol{c} \in C$. We begin by showing several useful properties of minimally $\mathbf{0}$-embracing sets and we present an algorithm that, given C, computes a minimally $\mathbf{0}$-embracing subset. Based on these results, we then show how to efficiently replace points in a $\mathbf{0}$-embracing simplex while preserving the embrace of the origin.

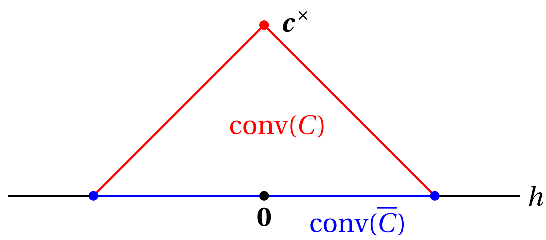

Figure 7.1.: The blue points constitute the linearly dependent set $\overline{C}$. The removal of $\boldsymbol{c}^\times$ maintains the embrace of the origin.

Lemma 7.1. *Let $C \subset \mathbb{R}^d$ be an affinely independent $\mathbf{0}$-embracing set. Then, a subset C' of C is linearly dependent if and only if C' embraces the origin.*

Proof. First, we observe that all $\mathbf{0}$-embracing subsets of C must be linearly dependent. Let now C' be a linearly dependent subset of C. We want to show that then C' is $\mathbf{0}$-embracing, too. Assume without loss of generality that C' is a proper subset and let $\boldsymbol{c}^\times \in C \setminus C'$ be a missing point. We prove that the set $\overline{C} = C \setminus \{\boldsymbol{c}^\times\}$ is $\mathbf{0}$-embracing. A repeated application of this argument then implies the statement.

Since $C' \subseteq \overline{C}$, the set $\overline{C}$ is linearly dependent. Thus, we can write $\mathbf{0}$ as a nontrivial linear combination $\sum_{\boldsymbol{c}\in\overline{C}} \phi_{\boldsymbol{c}}\boldsymbol{c}$ of the points in $\overline{C}$, where $\phi_{\boldsymbol{c}} \in \mathbb{R}$ for all $\boldsymbol{c} \in \overline{C}$. Furthermore, since C is affinely independent, so is $\overline{C}$, and hence $\sum_{\boldsymbol{c}\in\overline{C}} \phi_{\boldsymbol{c}} \neq 0$. By rescaling the coefficients, we obtain an affine combination of $\mathbf{0}$. This implies $\operatorname{aff}(\overline{C}) = \operatorname{span}(\overline{C})$. Now, because $\overline{C} = C \setminus \{\boldsymbol{c}^\times\}$ and because C is affinely independent, the point $\boldsymbol{c}^\times$ is not contained in the affine hull of $\overline{C}$ and thus not in the linear span of $\overline{C}$. Then, there exists a hyperplane h that contains $\operatorname{span}(\overline{C})$ but not $\boldsymbol{c}^\times$. See Figure 7.1. Because $\operatorname{conv}(C)$ is on one side of h, the intersection $h \cap \operatorname{conv}(C) = \operatorname{conv}(\overline{C})$ is a face of $\operatorname{conv}(C)$. Since h and $\operatorname{conv}(C)$ both contain the origin, the face $\operatorname{conv}(\overline{C})$ must contain the origin, too. Hence, $\overline{C}$ is $\mathbf{0}$-embracing. ■

Lemma 7.2. *Let $C \subset \mathbb{R}^d$ be a minimally $\mathbf{0}$-embracing set. Then, the following holds:*

(i) C is affinely independent and all proper subsets of C are linearly independent.

(ii) For all $\boldsymbol{c} \in C$, the point $-\boldsymbol{c}$ is ray-embraced by $C \setminus \{\boldsymbol{c}\}$.

In particular, $\dim C = |C| - 1$ *and* $\operatorname{pos}(C) = \operatorname{span}(C)$.

Proof. If C is affinely dependent, then by Theorem 1.1 there exists a proper subset that embraces the origin. Thus, C must be affinely independent. Hence, (i) is implied by Lemma 7.1. Write now C as $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_n$ and let $\lambda_1, \ldots, \lambda_n \in \mathbb{R}_+$ be coefficients that sum to 1 such that $\mathbf{0} = \sum_{i=1}^m \lambda_i \boldsymbol{c}_i$. Then, $-\lambda_i \boldsymbol{c}_i \in \operatorname{pos}(C)$ for all $i \in [n]$. Because $C \setminus \{\boldsymbol{c}\}$ does not embrace the origin for any $\boldsymbol{c} \in C$, we have $\lambda_i > 0$ for $i \in [n]$. This implies (ii). ■

Using the fact that all proper subsets of a minimally $\mathbf{0}$-embracing set C are linearly independent, we show how to compute for each point in the positive span of C the coefficients of the positive combination.

Lemma 7.3. *Let $C \subset \mathbb{R}^d$ be a minimally $\mathbf{0}$-embracing set and let $\boldsymbol{q} \in \operatorname{pos}(C)$ be a point. Then, we can compute the coefficients of the positive combination of $\boldsymbol{q}$ with the points in C in $O(d^4)$ time.*

Proof. First assume that $\boldsymbol{q} = \mathbf{0}$. Let $\boldsymbol{c} \in C$ be an arbitrary point and denote with $\overline{C} = C \setminus \{\boldsymbol{c}\}$ the remaining points. By Lemma 7.2, $-\boldsymbol{c}$ is ray-embraced by $\overline{C}$. Thus, the linear system $A\boldsymbol{x} = -\boldsymbol{c}$, where A is the matrix whose columns are the points from $\overline{C}$, has a solution. By Lemma 7.2 (i), the set $\overline{C}$ is linearly independent and hence this solution is unique. Thus, we can compute the coefficients $\psi_{\boldsymbol{c}} \in \mathbb{R}$, $\boldsymbol{c} \in \overline{C}$, such that $-\boldsymbol{c} = \sum_{\boldsymbol{c}\in\overline{C}} \psi_{\boldsymbol{c}}\boldsymbol{c}$ in $O(d^3)$ time with Gaussian elimination. Moreover, since the solution is unique, we must have $\psi_{\boldsymbol{c}} \geq 0$ for all $\boldsymbol{c} \in \overline{C}$. Set $\psi_{\boldsymbol{c}}$ to 1. Then, $\mathbf{0} = \sum_{\boldsymbol{c}\in C} \psi_{\boldsymbol{c}}\boldsymbol{c}$ and all coefficients are nonnegative.

Now assume that $\boldsymbol{q} \neq \mathbf{0}$. We iterate through all points in $\boldsymbol{c} \in C$ and solve the linear system $L_{\boldsymbol{c}} : A\boldsymbol{x} = \boldsymbol{q}$, where the columns of A are the points in $C \setminus \{\boldsymbol{c}\}$. Again by Lemma 7.2 (i), the columns of A are linearly independent and hence the solution $\boldsymbol{x}_{\boldsymbol{c}}$ to $L_{\boldsymbol{c}}$ is unique if it exists. If $\boldsymbol{x}_{\boldsymbol{c}} \geq \mathbf{0}$, we have found the desired coefficients. By Theorem 1.1, there exists a proper subset C' of C that ray-embraces $\boldsymbol{q}$ and thus there exists a point $\boldsymbol{c}^\times \in C$ for which $\boldsymbol{x}_{\boldsymbol{c}^\times} \geq \mathbf{0}$. Solving the linear system $L_{\boldsymbol{c}}$ takes $O(d^3)$ time for each point $\boldsymbol{c} \in C$ with Gaussian elimination, and hence we need $O(d^4)$ time in total before finding the $\boldsymbol{q}$-embracing subset $C \setminus \{\boldsymbol{c}^\times\}$ together with the coefficients of the positive combination. ■

We can now combine the previous results to show that given a $\mathbf{0}$-embracing set, we can find a minimally $\mathbf{0}$-embracing subset in polynomial time together with the coefficients of the convex combination of the origin.

Lemma 7.4. *Let $C \subset \mathbb{R}^d$ be a $\mathbf{0}$-embracing set of size n. Given the coefficients of the convex combination of $\mathbf{0}$ with the points in C, we can find a minimally $\mathbf{0}$-embracing subset $C' \subseteq C$ and the coefficients of the convex combination of $\mathbf{0}$ with the points in C' in $O\left(n^2 + nd^3 + d^4\right)$ time.*

Proof. First, we apply Theorem 1.1 to obtain an affinely independent subset C' of C that embraces the origin. Then, we iteratively test for each point $\boldsymbol{c} \in C'$ whether the set $C' \setminus \{\boldsymbol{c}\}$ is linearly dependent. If so, we remove $\boldsymbol{c}$ from C'. After iterating through all points, the resulting set still embraces the origin by Lemma 7.1 and moreover, since no proper subset is linearly dependent, it is minimally $\mathbf{0}$-embracing.

The initial application of Theorem 1.1 needs $O\left(nd^3 + n^2\right)$ time. Then, checking for one point $\boldsymbol{c} \in C'$ whether $C' \setminus \{\boldsymbol{c}\}$ is linearly dependent takes $O\left(d^3\right)$ time with Gaussian elimination. Because C' is affinely independent, we have $|C'| \leq d+1$ and thus the claimed running time follows. ■

Let now $Q \subset \mathbb{R}^d$ be a set and let $C \subset \mathbb{R}^d$ as before be a $\mathbf{0}$-embracing set. We say a subset C' of C is $\mathbf{0}$-*embracing equivalent* to Q with respect to C if $(C \setminus C') \cup Q$ embraces $\mathbf{0}$. In the following, we show that if Q embraces the origin when orthogonally projected onto $\operatorname{span}(C)^\perp$, there is always at least one point in C that is $\mathbf{0}$-embracing equivalent to Q. See Figure 7.2 (a).

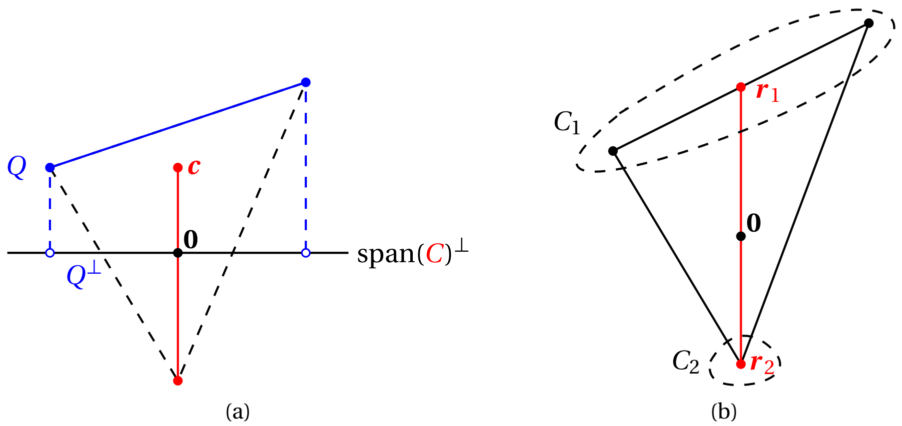

Figure 7.2.: (a) An example of Lemma 7.5. The red points constitute the minimal $\mathbf{0}$-embracing set C and the blue points constitute the set Q that embraces the origin when projected onto $\operatorname{span}(C)^\perp$. The point $\boldsymbol{c} \in C$ is $\mathbf{0}$-embracing equivalent to Q. (b) An example of Lemma 7.6. The set C consists of the vertices of the simplex, and the two representative points are with respect to the indicated partition.

Lemma 7.5. *Let $C \subset \mathbb{R}^d$ be a $\mathbf{0}$-embracing set and let Q be a set whose orthogonal projection $Q^\perp$ onto* $\operatorname{span}(C)^\perp$ *embraces $\mathbf{0}$. Then, there exists a point $\boldsymbol{c} \in C$ that is $\mathbf{0}$-embracing equivalent to Q*

with respect to C. *Furthermore, if both* C *and* $Q^\perp$ *are minimally* $\mathbf{0}$*-embracing, we can compute* $\boldsymbol{c}$ *together with the coefficients of the convex combination of* $\mathbf{0}$ *with the points in* $Q \cup C \setminus \{\boldsymbol{c}\}$ *in* $O(d^4)$ *time.*

Proof. We first prove that there is always a point in C that is $\mathbf{0}$-embracing equivalent to Q and then show how to find this point efficiently. We can assume without loss of generality that C is minimally $\mathbf{0}$-embracing since otherwise the statement holds trivially. Let now $\boldsymbol{q}_1, \dots, \boldsymbol{q}_m \in \mathbb{R}^d$ denote the points in Q and write each $\boldsymbol{q}_i$, $i \in [m]$, as the sum of a vector $\boldsymbol{p}_i \in \operatorname{span}(C)$ and a vector $\boldsymbol{p}_i^\perp \in \operatorname{span}(C)^\perp$. Because Q projected onto $\operatorname{span}(C)^\perp$ is $\mathbf{0}$-embracing, there are coefficients $\lambda_1, \dots, \lambda_m \in \mathbb{R}_+$ that sum to 1 such that $\mathbf{0} = \sum_{i=1}^m \boldsymbol{p}_i^\perp$. Consider the convex combination $\boldsymbol{q} = \sum_{i=1}^m \lambda_i \boldsymbol{q}_i$ of the points in Q with the same coefficients. Since

$$\boldsymbol{q} = \sum_{i=1}^m \lambda_i \left(\boldsymbol{p}_i + \boldsymbol{p}_i^\perp\right) = \left(\sum_{i=1}^m \lambda_i \boldsymbol{p}_i\right) + \left(\sum_{i=1}^m \lambda_i \boldsymbol{p}_i^\perp\right) = \sum_{i=1}^m \lambda_i \boldsymbol{p}_i,$$

the point $\boldsymbol{q}$ is contained in $\operatorname{span}(C)$. By Lemma 7.2, we have $\operatorname{pos}(C) = \operatorname{span}(C)$ and hence $-\boldsymbol{q}$ is ray-embraced by C. Now, the cone version of Theorem 1.1 states that there is a linearly independent subset C' of C that ray-embraces $-\boldsymbol{q}$. Because $\dim C = |C| - 1$ by Lemma 7.2, the set C' must be a proper subset. Then, Q is $\mathbf{0}$-embracing equivalent to all points in $C \setminus C' \neq \emptyset$.

It remains to show how to find a point in $C \setminus C'$. We assume that both C and $Q^\perp$ are minimally $\mathbf{0}$-embracing, where $Q^\perp$ is the orthogonal projection of Q onto $\operatorname{span}(C)^\perp$. Using the algorithm from Lemma 7.3, we compute the coefficients of the convex combination of the origin with the points in $Q^\perp$ and hence the point $-\boldsymbol{q}$ in $O(d^4)$ time. Applying Lemma 7.3 again, we can determine the coefficients of the positive combination of $-\boldsymbol{q}$ with the points in C in $O(d^4)$ time. We use the algorithm from Lemma 7.4 to find a minimally $(-\boldsymbol{q})$-embracing subset C' from C in $O(d^4)$ time. Then, we can choose any point in $C \setminus C'$ as $\boldsymbol{c}$. Finally, since we know the coefficients of the convex combination of $\boldsymbol{q}$ with the points in Q and since we can apply Lemma 7.3 to compute the coefficients of the positive combination of $-\boldsymbol{q}$ with the points in C', we can compute the coefficients of the convex combination of the origin with the points in $C' \cup Q$ by rescaling appropriately. The algorithm takes in total $O(d^4)$ time, as claimed. ■

Lemma 7.5 itself does not yet lead to a nontrivial approximation algorithm due to the weak guarantee that only a single point in C is $\mathbf{0}$-embracing equivalent to Q. To amplify the number of points that can be replaced, we conclude this section by showing how to compute a set of *representative points* R for C that each stand for a specific subset of C such that if a point in R is $\mathbf{0}$-embracing equivalent to a set Q with respect to R, then the corresponding subset of C is $\mathbf{0}$-embracing equivalent to Q with respect to C, too. See Figure 7.2 (b).

Lemma 7.6. *Let* $C \subset \mathbb{R}^d$ *be a minimally* $\mathbf{0}$*-embracing set and let* $C_1, \dots, C_m$ *be a partition of* C *into* $m \geq 2$ *sets with* $|C_i| \geq 1$ *for all* $i \in [m]$. *Then, we can compute in* $O(d^4)$ *time a set of points* $R = \{\boldsymbol{r}_1, \dots, \boldsymbol{r}_m\} \subset \mathbb{R}^d$ *with the following properties:*

(i) R *is minimally* $\mathbf{0}$*-embracing.*

(ii) Let $Q \subset \mathbb{R}^d$ *be a set that is* $\mathbf{0}$*-embracing equivalent to some point* $\boldsymbol{r}_j \in R$ *with respect to* R. *Then,* Q *is* $\mathbf{0}$*-embracing equivalent to* C_j *with respect to* C.

We call the points in R representative points *for C with respect to the partition $C_1, \dots, C_m$.*

Proof. Since C is minimally $\mathbf{0}$-embracing, we can write $\mathbf{0}$ as a convex combination $\sum_{\boldsymbol{c} \in C} \lambda_{\boldsymbol{c}} \boldsymbol{c}$ such that all $\lambda_{\boldsymbol{c}}$ are strictly greater than 0 and sum to 1. With the algorithm from Lemma 7.3, we can compute these coefficients in $O\left(d^4\right)$ time. For $i \in [m]$, set $\boldsymbol{r}_i$ to $\sum_{\boldsymbol{c} \in C_i} \lambda_{\boldsymbol{c}} \boldsymbol{c}$. Clearly, R is $\mathbf{0}$-embracing. Moreover, for all $j \in [m]$, the set $\left\{\boldsymbol{r}_i \,\middle|\, i \in [m],\, i \neq j\right\}$ is not $\mathbf{0}$-embracing since otherwise the set $\bigcup_{i=1,\, i\neq j}^{m} C_i$, a strict subset of C, is $\mathbf{0}$-embracing, a contradiction to C being minimally $\mathbf{0}$-embracing. Let now Q be a set that is $\mathbf{0}$-embracing equivalent to some point $\boldsymbol{r}_j \in R$ with respect to R. That is, the set $Q \cup \left(R \setminus \{\boldsymbol{r}_j\}\right)$ embraces the origin. Because $\boldsymbol{r}_i \in \operatorname{pos}(C_i)$ for $i \in [m]$, then the set $Q \cup \left(\bigcup_{i=1,\, i\neq j}^{m} C_i\right)$ is $\mathbf{0}$-embracing as well, and hence Q is $\mathbf{0}$-embracing equivalent to C_j with respect to C. ■

7.2. Computing k-Colorful Choices

Lemmas 7.5 and 7.6 suggest a simple approximation algorithm. Let $C_1, \dots, C_m \subset \mathbb{R}^d$ be m color classes that each embrace the origin for $i \in [m]$ and set $k = \max\left(d - m + 2, \left\lceil \frac{d+1}{2} \right\rceil\right)$. Then, the following algorithm recursively computes a $\mathbf{0}$-embracing k-colorful choice. First, we prune C_1 with Lemma 7.4 and partition it into two sets C_1', C_2' of size at most $\lceil (d+1)/2 \rceil$. Using Lemma 7.6, we compute two representatives points $\boldsymbol{r}_1$, $\boldsymbol{r}_2$ for this partition of C_1. Then, we project the remaining $m-1$ color classes onto the $(d-1)$-dimensional space that is orthogonal to $\operatorname{span}(\boldsymbol{r}_1, \boldsymbol{r}_2)^{\perp}$ and we recursively compute a $\mathbf{0}$-embracing k-colorful choice Q with respect to the projections of $C_2, \dots, C_m$. By Lemmas 7.5 and 7.6, one of the two sets C_1', C_2', say C_1', is $\mathbf{0}$-embracing equivalent to Q with respect to C_1. Since Q is a k-colorful choice that does not contain points from C_1 and since $|C_1'|, |C_2'| \le k$, the set $C_1' \cup Q$ is a $\mathbf{0}$-embracing k-colorful choice. The recursion stops once only one color class is left. Then, we are in dimension $d - m + 1$. Since $d - m + 2 \le k$, pruning the single remaining color class with Lemma 7.4 results already in a $\mathbf{0}$-embracing k-colorful choice. For details, see Algorithm 7.1.

Algorithm 7.1: Simple Approximation

Input: m sets $C_1, \dots, C_m \subset \mathbb{R}^d$ that each embrace the origin, and for each C_i, $i \in [m]$, the coefficients of the convex combination of $\mathbf{0}$ with the points in C_i

Output: $\mathbf{0}$-embracing $\max\left(d - m + 2, \left\lceil \frac{d+1}{2} \right\rceil\right)$-colorful choice

$C \leftarrow$ prune C_1 with Lemma 7.4;
if $m = 1$ **then return** C;
$C_1', C_2' \leftarrow$ partition of C into two sets, each of size at most $\left\lceil \frac{d+1}{2} \right\rceil$;
Compute representative points $\boldsymbol{r}_1$, $\boldsymbol{r}_2$ for C_1', C_2';
$\check{C}_2, \dots, \check{C}_m \leftarrow$ orthogonal projection of $C_2, \dots, C_m$ onto $\operatorname{span}(\boldsymbol{r}_1, \boldsymbol{r}_2)^{\perp}$;
$\check{Q} \leftarrow \texttt{recurse}(\check{C}_2, \dots, \check{C}_m)$;
$Q \leftarrow$ replace projected points in Q by original points from $\bigcup_{i=2}^{m} C_i$;
Determine which point $\boldsymbol{r}_{i^\times} \in \{\boldsymbol{r}_1, \boldsymbol{r}_2\}$ is $\mathbf{0}$-embracing equivalent to Q;
return $\left(C \setminus C_{i^\times}'\right) \cup Q$ *pruned with Lemma 7.4*;

Theorem 7.7. *Let $C_1, \dots, C_m \subset \mathbb{R}^d$ be $m \le d$ color classes such that C_i is a $\mathbf{0}$-embracing set of size $O(d)$ for $i \in [k]$. On input $C_1, \dots, C_m$ and given the coefficients of the convex combination of the origin for each set C_i, Algorithm 7.1 computes a $\mathbf{0}$-embracing $\max\left(d - m + 2, \left\lceil \frac{d+1}{2} \right\rceil\right)$-colorful choice in $O\left(d^5\right)$ time. In particular, for $m = \lfloor d/2 \rfloor + 1$, the algorithm computes a $(\lceil d/2 \rceil + 1)$-colorful choice.*

Proof. The correctness of Algorithm 7.1 is a direct consequence of Lemmas 7.5 and 7.6. It remains to analyze the running time. In each step of the recursion except for the last one, we prune two times a set of size $O(d)$ with Lemma 7.4. This needs $O\left(d^4\right)$ time. Furthermore, by Lemma 7.6, computing two representative points takes $O\left(d^4\right)$ time, too. Finally, given the set Q, determining which representative point is $\mathbf{0}$-embracing equivalent to Q takes also $O\left(d^4\right)$ by Lemma 7.5. Thus, we need $O\left(d^4\right)$ time per step of the recursion and there are $O(d)$ recursion steps in total. Hence, the total running time is $O\left(d^5\right)$. ∎

Although nontrivial, the fact that we can take in polynomial time half of the points from each color class to construct a $\mathbf{0}$-embracing $(\lceil d/2 \rceil + 1)$-colorful choice may not be too surprising. In the remainder of this chapter, we present a generalization of Algorithm 7.1 that computes $\mathbf{0}$-embracing $\lceil \varepsilon d \rceil$-colorful choices in polynomial time for any fixed $\varepsilon > 0$. The improved approximation guarantee is achieved by repeatedly replacing subsets of C with Lemmas 7.5 and 7.6 in each step of the recursion. To still ensure polynomial running time, we reduce the dimensionality by a constant fraction in each step of the recursion. Additionally, we slightly worsen the approximation guarantee in each level of the recursion, i.e., if the current recursion level is j and the dimensionality is d', then we do not compute an $\lceil \varepsilon d' \rceil$-colorful choice but a $\left\lceil (1-\varepsilon/2)^{-j/2} \varepsilon d' \right\rceil$-colorful choice. As we will see, this additional "slack" in the approximation guarantee limits the recursion depth to a constant depending only on ε as after only $O(1)$ recursion steps even a complete color class pruned with Lemma 7.4 is already a good enough approximation.

In more detail, let $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ be $d+1$ sets that each embrace the origin and let $\varepsilon > 0$ be a parameter. We want to compute an $\lceil \varepsilon d \rceil$-colorful choice that embraces the origin. Set

$$d_j = \left\lceil \left(1 - \frac{\varepsilon}{2}\right)^j d \right\rceil \text{ and } k_j = \left\lceil \varepsilon \left(1 - \frac{\varepsilon}{2}\right)^{j/2} d \right\rceil$$

for $j \in \mathbb{N}$. The sequence d_j controls the dimension reduction argument with Lemmas 7.5 and 7.6, i.e., in the jth recursion level, the dimensionality of the input will be d_j. The sequence k_j defines the approximation guarantee in the jth recursion level. Note that $d_0 = d$ and $k_0 = \lceil \varepsilon d \rceil$. Assume now we are in recursion level j. That is, the input consists of $d_j + 1$ color classes $C_1, \dots, C_{d_j+1} \subset \mathbb{R}^{d_j}$ that each embrace the origin together with the coefficients of their convex combinations of the origin and we want to compute a $\mathbf{0}$-embracing k_j-colorful choice. As in the previous algorithm, we begin by computing a minimal $\mathbf{0}$-embracing subset C of C_1 with Lemma 7.4. If $k_j \ge d_j + 1$, then C is already a valid approximation. Otherwise, we repeatedly replace subsets of C until it contains at most k_j points from each color as follows. Set $m = d_j - d_{j+1} + 1$. We partition C into sets $C'_1, \dots, C'_m$ by distributing the points from each color in C equally among these m sets and we compute representative points $\boldsymbol{r}_1, \dots, \boldsymbol{r}_m$ for this partition. Let $C^\star_1, \dots, C^\star_{d_{j+1}+1} \in \left\{C_2, \dots, C_{d_j+1}\right\}$ be $d_{j+1} + 1$ color classes, where we discuss shortly how they are chosen. Then, we recursively compute a k_{j+1}-colorful choice Q for $C_1^\star, \dots, C_{d_{j+1}+1}^\star$ that embraces the origin when projected on $U = \operatorname{span}(\boldsymbol{r}_1, \dots, \boldsymbol{r}_m)^\perp$. Note that $\dim U = d_j - (m-1) = d_{j+1}$ and hence the dimensionality of the input in recursion level $j+1$ is d_{j+1}, as desired. Then, by Lemmas 7.5 and 7.6, at least one reference point $\boldsymbol{r}_{i^\times}$ and hence at least one of the sets $C'_{i^\times}$ is $\mathbf{0}$-embracing equivalent to Q. We set C to $(C \setminus C_{i^\times}) \cup Q$ and prune it with Lemma 7.4. We repeat these steps until C is a k_j-colorful choice. To ensure progress, m should be smaller than k_j so that $C_{i^\times}$ is guaranteed to contain a point from each color that appears more than k_j times in C. Furthermore, Q should not contain points with colors that appear "often" in C. We call a color class C_i *light* with respect to C if $|C \cap C_i| \le k_j - k_{j+1}$ and otherwise *heavy*. For the recursion, we only use light color classes. A k_{j+1}-colorful choice with light colors can be added safely to C without increasing any color over the threshold k_j. In particular, since we start with $C = C_1$ and only use light color classes, no other color class can ever occur more than k_j times in C and hence we are finished once the number of points from C_1 is at most k_j. Please refer to Algorithm 7.2 for details.

Algorithm 7.2: $\lceil \varepsilon d \rceil$-Approximation

Input: recursion depth $j \in \mathbb{N}_0$ (initially 0), original dimension $d \in \mathbb{N}$, approximation parameter $\varepsilon > 0$, $d_j + 1$ sets $C_1, \dots, C_{d_j+1} \subset \mathbb{R}^{d_j}$ that each embrace the origin, and for each C_i the coefficients of the convex combination of $\mathbf{0}$ with the points in C_i

Output: $\mathbf{0}$-embracing k_j-colorful choice

$k_j \leftarrow \left\lceil \varepsilon \left(1 - \frac{\varepsilon}{2}\right)^j d \right\rceil$;
$d_{j+1} \leftarrow \left\lceil \left(1 - \frac{\varepsilon}{2}\right)^{j+1} d \right\rceil$;
$m \leftarrow d_j - d_{j+1} + 1$;
$C \leftarrow$ prune C_1 with Lemma 7.4;
while $|C \cap C_1| > k_j$ **do**
  $C'_1, \dots, C'_m \leftarrow$ partition of C s.t. the points from each color class are evenly distributed;
  Compute representative points $\boldsymbol{r}_1, \dots, \boldsymbol{r}_m$ for $C'_1, \dots, C'_m$ with Lemma 7.6;
  Find $d_{j+1}+1$ light color classes $C_1^\star, \dots, C_{d_{j+1}+1}^\star \in \left\{C_2, \dots, C_{d_j+1}\right\}$;
  $\breve{C}_1, \dots, \breve{C}_{d_{j+1}+1} \leftarrow$ orthogonal projection of $C_1^\star, \dots, C_{d_{j+1}+1}^\star$ onto $\operatorname{span}(\boldsymbol{r}_1, \dots, \boldsymbol{r}_m)^\perp$;
  $\breve{Q} \leftarrow$ `recurse(`$j+1, d, \varepsilon, \breve{C}_1, \dots, \breve{C}_{d_{j+1}+1}$`)`;
  $Q \leftarrow$ replace projected points in Q by original points from $\bigcup_{i=1}^{d_{j+1}+1} C_i^\star$;
  Determine which point $\boldsymbol{r}_{i^\times} \in \{\boldsymbol{r}_1, \dots, \boldsymbol{r}_m\}$ is $\mathbf{0}$-embracing equivalent to Q with Lemma 7.5;
  $C \leftarrow \left(C \setminus C'_{i^\times}\right) \cup Q$ pruned with Lemma 7.4;
return C;

The next lemma states that for ε fixed, the number of necessary recursions before a trivial approximation with Lemma 7.4 suffices is constant.

Lemma 7.8. *For any $\varepsilon = \Omega\left(d^{-1}\right)$ there exists a $j = \Theta\left(\varepsilon^{-1} \ln \varepsilon^{-1}\right)$ such that $k_j \ge d_j + 1$.*

Proof. Replacing d_j with its definition, we obtain

$$d_j + 1 = \left\lceil \left(1 - \frac{\varepsilon}{2}\right)^j d \right\rceil + 1 \leq \left(1 - \frac{\varepsilon}{2}\right)^j d + 2. \tag{7.1}$$

Using $-\ln\left(1 - \frac{\varepsilon}{2}\right) \leq \varepsilon$ if $\varepsilon \leq 1$, we have for $j = O\left(\varepsilon^{-1} \ln d\right)$

$$\left(1 - \frac{\varepsilon}{2}\right)^j d \geq 1. \tag{7.2}$$

Furthermore, using the fact that $-\ln\left(1 - \frac{\varepsilon}{2}\right) \geq \frac{\varepsilon}{2}$, we have for $j = \Omega\left(\varepsilon^{-1} \ln \varepsilon^{-1}\right)$

$$3\left(1 - \frac{\varepsilon}{2}\right)^{j/2} \leq \varepsilon. \tag{7.3}$$

Combining (7.2) and (7.3) with (7.1), we get

$$d_j + 1 \leq 3\left(1 - \frac{\varepsilon}{2}\right)^j d \leq \varepsilon\left(1 - \frac{\varepsilon}{2}\right)^{j/2} d \leq \left\lceil \varepsilon\left(1 - \frac{\varepsilon}{2}\right)^{j/2} d \right\rceil = k_j$$

and hence, if $d = \Omega\left(\varepsilon^{-1}\right)$ and $j = \Theta\left(\varepsilon^{-1} \ln \varepsilon^{-1}\right)$, we have $d_j + 1 \leq k_j$. ■

Next, we show that if the recursion depth is not too large, then we can always find enough light color classes.

Lemma 7.9. *Let $j \in \mathbb{N}$ and let $C_1, \ldots, C_{d_j+1} \subset \mathbb{R}^{d_j}$ be $d_j + 1$ color classes. Furthermore, let $C \subseteq \bigcup_{i=1}^{d_j+1} C_i$ be a set of size at most $d_j + 1$. For all $j = O\left(\varepsilon^{-1} \ln \varepsilon^3 d\right)$, there exist $d_{j+1} + 1$ light color classes with respect to C.*

Proof. We recall that a color class C_i, $i \in [d_j + 1]$, is light with respect to C if $|C \cap C_i| \leq k_j - k_{j+1}$. Then, the number of heavy color classes h is bounded by

$$h \leq \left\lceil \frac{d_j + 1}{k_j - k_{j+1}} \right\rceil \leq \frac{2d_j}{k_j - k_{j+1}} + 1, \tag{7.4}$$

since $d_j \geq 1$ for all $j \in \mathbb{N}$. We can bound the denominator as follows

$$\begin{aligned} k_j - k_{j+1} &= \left\lceil \varepsilon\left(1 - \frac{\varepsilon}{2}\right)^{j/2} d \right\rceil - \left\lceil \varepsilon\left(1 - \frac{\varepsilon}{2}\right)^{(j+1)/2} d \right\rceil \\ &\geq \varepsilon\left(1 - \frac{\varepsilon}{2}\right)^{j/2} d - \varepsilon\left(1 - \frac{\varepsilon}{2}\right)^{(j+1)/2} d - 1 \\ &= \varepsilon\left(1 - \frac{\varepsilon}{2}\right)^{j/2} d\left(1 - \sqrt{1 - \frac{\varepsilon}{2}}\right) - 1 \geq \frac{\varepsilon^2}{4}\left(1 - \frac{\varepsilon}{2}\right)^{j/2} d - 1, \end{aligned} \tag{7.5}$$

where we apply $1 - \sqrt{1 - \frac{\varepsilon}{2}} \geq \frac{\varepsilon}{4}$ in the last inequality. Using that $-\ln\left(1 - \frac{\varepsilon}{2}\right) \leq \varepsilon$ if $\varepsilon \leq 1$, we have for $j = O\left(\varepsilon^{-1} \ln \varepsilon^2 d\right)$

$$\frac{\varepsilon^2}{4}\left(1 - \frac{\varepsilon}{2}\right)^{j/2} d \geq 2 \tag{7.6}$$

and hence (7.5) can be simplified to

$$k_j - k_{j+1} \geq \frac{\varepsilon^2}{8}\left(1-\frac{\varepsilon}{2}\right)^{j/2} d. \tag{7.7}$$

Plugging (7.7) into (7.4) and using (7.6), we obtain

$$h \leq \frac{2\left\lceil\left(1-\frac{\varepsilon}{2}\right)^j d\right\rceil}{\frac{\varepsilon^2}{8}\left(1-\frac{\varepsilon}{2}\right)^{j/2} d} + 1 \leq \frac{2\left(1-\frac{\varepsilon}{2}\right)^j d}{\frac{\varepsilon^2}{8}\left(1-\frac{\varepsilon}{2}\right)^{j/2} d} + 3 = \frac{16}{\varepsilon^2}\left(1-\frac{\varepsilon}{2}\right)^{j/2} + 3.$$

Then, the number of light color classes ℓ is at least

$$\begin{aligned}\ell = d_j + 1 - h &\geq \left\lceil\left(1-\frac{\varepsilon}{2}\right)^j d\right\rceil - \frac{16}{\varepsilon^2}\left(1-\frac{\varepsilon}{2}\right)^{j/2} - 2 \\ &\geq \left(1-\frac{\varepsilon}{2}\right)^j d\left(1 - \frac{16}{\varepsilon^2\left(1-\frac{\varepsilon}{2}\right)^{j/2} d} - \frac{2}{\left(1-\frac{\varepsilon}{2}\right)^j d}\right). \end{aligned} \tag{7.8}$$

For $j = O\left(\varepsilon^{-1}\ln \varepsilon^3 d\right)$, we have

$$\frac{16}{\varepsilon^2\left(1-\frac{\varepsilon}{2}\right)^{j/2} d} + \frac{2}{\left(1-\frac{\varepsilon}{2}\right)^j d} \leq \frac{\varepsilon}{4}$$

and thus (7.8) implies

$$\ell \geq \left(1-\frac{\varepsilon}{4}\right)\left(1-\frac{\varepsilon}{2}\right)^j d. \tag{7.9}$$

For $j = O\left(\varepsilon^{-1}\ln \varepsilon d\right)$, we can bound

$$\frac{\varepsilon}{4}\left(1-\frac{\varepsilon}{2}\right)^j d \geq 2. \tag{7.10}$$

Combining (7.10) with (7.9), we get

$$\ell \geq \left(1-\frac{\varepsilon}{4}\right)\left(1-\frac{\varepsilon}{2}\right)^j d + \left(\frac{\varepsilon}{4}-\frac{\varepsilon}{4}\right)\left(1-\frac{\varepsilon}{2}\right)^j d \geq \left(1-\frac{\varepsilon}{2}\right)^{j+1} d + 2 \geq \left\lceil\left(1-\frac{\varepsilon}{2}\right)^{j+1} d\right\rceil + 1 = d_{j+1} + 1.$$

Thus, for $j = O\left(\varepsilon^{-1}\ln \varepsilon^3 d\right)$, there are at least $d_{j+1}+1$ light color classes with respect to C. ■

Before we finally prove correctness, we show if the recursion depth j is not too large, then each set of the partition of C contains at least one point from C_1 until C is a k_j-colorful choice. This implies that each iteration of the while-loop decreases the amount of points from C_1 in C.

Lemma 7.10. *For all $j = O\left(\varepsilon^{-1}\ln \varepsilon d\right)$, we have $m = d_j - d_{j+1} + 1 \leq k_j + 1$.*

Proof. First, we upper bound m as follows:

$$\begin{aligned} m = d_j - d_{j+1} + 1 &= \left\lceil \left(1-\frac{\varepsilon}{2}\right)^j d \right\rceil - \left\lceil \left(1-\frac{\varepsilon}{2}\right)^{j+1} d \right\rceil + 1 \\ &\leq \left(1-\frac{\varepsilon}{2}\right)^j d - \left(1-\frac{\varepsilon}{2}\right)^{j+1} d + 2 = \frac{\varepsilon}{2}\left(1-\frac{\varepsilon}{2}\right)^j d + 2. \end{aligned} \tag{7.11}$$

For $j = O\left(\varepsilon^{-1} \ln \varepsilon d\right)$, we can lower bound $\frac{\varepsilon}{2}\left(1-\frac{\varepsilon}{2}\right)^j d$ by 1. Using this in (7.11), we get

$$m \leq \varepsilon\left(1-\frac{\varepsilon}{2}\right)^j d + 1 \leq \left\lceil \varepsilon\left(1-\frac{\varepsilon}{2}\right)^j d \right\rceil + 1 = k_j + 1,$$

as desired. ■

Theorem 7.11. *Let $C_1, \ldots, C_{d+1} \subset \mathbb{R}^d$ be $d+1$ sets such that C_i is a $\mathbf{0}$-embracing set of size $O(d)$ for $i \in [d+1]$ and let $\varepsilon = \Omega\left(d^{-1/4}\right)$ be a parameter. On input 0, d, ε, $C_1, \ldots, C_{d+1}$, and given the coefficients of the convex combination of the origin with the points in C_i for $i \in [d+1]$, Algorithm 7.2 computes a $\mathbf{0}$-embracing $\lceil \varepsilon d \rceil$-colorful choice in $d^{O(\varepsilon^{-1} \ln \varepsilon^{-1})}$ time.*

Proof. We begin by showing that if the algorithm enters the while loop in recursion level j, it is always possible to find $d_{j+1}+1$ light color classes and that the projections $\breve{C}_1, \ldots, \breve{C}_{d_{j+1}+1}$ of these color classes are $\mathbf{0}$-embracing subsets of $\mathbb{R}^{d_{j+1}}$ (Line 9). In other words, we show that recursion is possible if C is not a k_j-colorful choice. Assume now the algorithm enters the while loop in recursion level j. Then, C is a minimally $\mathbf{0}$-embracing subset of $C_1 \subset \mathbb{R}^{d_j}$ and has size at least $k_j + 1$. In Line 6, we partition C into m sets $C'_1, \ldots, C'_m$ by distributing the points from each color class equally. By Lemma 7.10, we have $m \leq k_j + 1$ for $j = O\left(\varepsilon^{-1} \ln \varepsilon d\right)$, and hence each set C'_i is nonempty. Thus, the algorithm from Lemma 7.6 can be applied in Line 7 to compute the representative points $\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m$. Moreover $\dim \operatorname{span}(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m) = m-1$ by Lemma 7.6 and Lemma 7.2. Thus, $\dim \operatorname{span}(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m)^\perp = d - m + 1 = d_{j+1}$. Now, Lemma 7.9 guarantees that we can always find $d_{j+1}+1$ light color classes $C^\star_1, \ldots, C^\star_{d_{j+1}+1}$ if $j = O\left(\varepsilon^{-1} \ln \varepsilon^3 d\right)$. Because each color class $C^\star_i$, $i \in [d_{j+1}+1]$, is $\mathbf{0}$-embracing, so are their orthogonal projections onto $\operatorname{span}(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_k)^T$. Thus, recursion is possible if $j = O\left(\varepsilon^{-1} \ln \varepsilon^3 d\right)$. By Lemma 7.8, the recursion depth is limited to $\Theta\left(\varepsilon^{-1} \ln \varepsilon^{-1}\right)$ since then pruning C_1 with Lemma 7.4 in Line 4 is already a $\mathbf{0}$-embracing k_j-colorful choice. In this case, the while loop is never executed. We conclude that for $\varepsilon = O\left(d^{-1/4}\right)$, recursion is always possible as long as C is not a k_j-colorful choice.

Next, we prove that the algorithm computes in recursion level j a $\mathbf{0}$-embracing k_j-colorful choice. As discussed above, the recursion terminates after $O\left(\varepsilon^{-1} \ln \varepsilon^{-1}\right)$ steps when the set C from Line 4 is already a $\mathbf{0}$-embracing k_j-colorful choice. If C is not already a valid approximation, the while loop is executed. In each iteration of the while loop, C is partitioned into m sets $C'_1, \ldots, C'_m$ by distributing the points from each color equally among the C'_i. By Lemma 7.10, $m \leq k_j + 1$ for $j = O\left(\varepsilon^{-1} \ln \varepsilon d\right)$ and hence each set C'_i, $i \in [m]$, contains at least one point from C_1. Applying Lemmas 7.5 and 7.6, one of these sets, say $C'_{i^\times}$, is replaced in C by a recursively computed k_{j+1}-colorful choice Q that is $\mathbf{0}$-embracing when projected onto $\operatorname{span}(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m)^\perp$. Since we use in the recursion only light color classes with respect to C, and since C_1 is not a light color class, each iteration of the while loop strictly decreases the number of points from C_1 in C. Moreover, because Q contains only points from light color classes and since it is a k_{j+1}-colorful choice, $(C \setminus C_{i^\times}) \cup Q$ contains at most k_j points from the color classes $C_2, \ldots, C_{d_j+1}$. Thus, after $O(d)$ iterations, C is a $\mathbf{0}$-embracing k_j-colorful choice.

It remains to analyze the running time. The initial computation of C in Line 4 and each iteration of the while loop except for the recursive call takes $O\left(d^4\right)$ time. Since the while loop is executed $O(d)$ times and since the recursion depth is bounded by $O\left(\varepsilon^{-1}\ln\varepsilon^{-1}\right)$, the total running time of Algorithm 7.2 is $d^{O(\varepsilon^{-1}\ln\varepsilon^{-1})}$. ∎

7.3. Applications

Our main motivation to study algorithms for k-colorful choices was their potential application to approximate the descendants of ColorfulCarathéodory. We now give precise bounds on the quality and the running time of approximation algorithms that we obtain by combining algorithms for k-colorful choices with the presented reductions to ColorfulCarathéodory from Chapter 1. Unfortunately, the approximation guarantee of Algorithm 7.2 is too weak to obtain a nontrivial approximation algorithm for Tverberg and therefore also for Centerpoint and for SimplicialCenter. On the positive side, it leads to a nontrivial approximation algorithm for ColorfulKirchberger.

In the following, let $\mathcal{A}$ be an algorithm that given $d+1$ color classes $C_1, \ldots, C_{d+1} \subset \mathbb{R}^d$, each embracing the origin and of size $O(d)$, and for each C_i the coefficients of the convex combination of the origin, outputs a $\mathbf{0}$-embracing $k(d)$-colorful choice in $T(d)$ time, where $k, T : \mathbb{N} \mapsto \mathbb{N}$ are arbitrary but fixed functions.

Corollary 7.12. *Let $P \subset \mathbb{R}^d$ be a point set of size n and let $\mathcal{A}$ be as above. Set*

$$\widetilde{m} = \left\lceil \frac{n}{(d+1)^2\big(k(n-1)-1\big)+d+1} \right\rceil = \Omega\left(\frac{n}{d^2 k(n-1)}\right).$$

Then, a Tverberg $\widetilde{m}$-partition $\mathcal{T}$ of P and a point $\boldsymbol{p} \in\in \bigcap_{T\in\mathcal{T}} \operatorname{conv}(T)$ can be computed in $O\big((d^2+m)n^2 + T(n-1)\big)$ time.

Proof. Set $m = \lceil n/(d+1)\rceil$. In the proof of Theorem 1.8 on Page 12, we lift m copies of P with Lemma 1.9 to $\mathbb{R}^{n-1}$. Lifting one point needs $O(dm) = O(n)$ time and hence lifting all m copies takes $O\left(mn^2\right)$ time. Then, each point $\boldsymbol{p}_i \in \mathbb{R}^d$ from P corresponds to a color class $C_i = \left\{\hat{\boldsymbol{p}}_{i,j} \,\middle|\, j \in [m]\right\} \subset \mathbb{R}^{n-1}$ of size m and a $\mathbf{0}$-embracing colorful choice of $C_1, \ldots, C_n$ corresponds to the Tverberg partition $\mathcal{T} = \{T_1, \ldots, T_m\}$ that we obtain by assigning $\boldsymbol{p}_i \in P$ to T_j if $\hat{\boldsymbol{p}}_{i,j} \in C$. By construction of the color classes in the proof of Theorem 1.8, the barycenter of C_i is the origin for $i \in [n]$. Since we know then for each color class the coefficients of the convex combination of the origin, we can apply $\mathcal{A}$ to obtain a $\mathbf{0}$-embracing $k(n-1)$-colorful choice $\widetilde{C} \subseteq \bigcup_{i=1}^n C_i$ together with the coefficients of the convex combination of the origin with the points in $\widetilde{C}$. Let $\widetilde{T} = \left\{\widetilde{T}_1, \ldots, \widetilde{T}_m\right\}$ be a family of subsets of P that we construct as before by assigning $\boldsymbol{p}_i$ to $\widetilde{T}_j$ if $\boldsymbol{p}_{i,j} \in \widetilde{C}$. Here, $\widetilde{T}$ is a multiset, i.e., we allow $\widetilde{T}_i = \widetilde{T}_j$ for $i \neq j$. Since $\widetilde{C}$ embraces the origin, Lemma 1.9 guarantees that the intersection $\bigcap_{i=1}^m \operatorname{conv}\left(\widetilde{T}_i\right)$ is nonempty. Moreover, because we know the coefficients of the convex combination of the origin with the points in $\widetilde{C}$, we can compute in $O(dn)$ time a point $\boldsymbol{p}^\star \in \bigcap_{i=1}^m \operatorname{conv}(\widetilde{T}_i)$ together with the coefficients of the convex combination of $\boldsymbol{p}^\star$ with the points in $\widetilde{T}_i$ for $i \in [m]$, as described in the proof of Lemma 1.9.

Now, we construct a Tverberg partition for P out of $\widetilde{\mathcal{T}}$ by a greedy strategy that iteratively removes sets from $\widetilde{\mathcal{T}}$. Let $\widetilde{T} \in \widetilde{\mathcal{T}}$ be some set and remove it from $\widetilde{\mathcal{T}}$. Since we know the coefficients of the convex combination of $\boldsymbol{p}^\star$ with the points in $\widetilde{T}$, Theorem 1.1 can be applied to prune $\widetilde{T}$ to a $\boldsymbol{p}^\star$-embracing set of size at most $d+1$ in $O(d^3n+n^2)$ time. Then, for each point $\boldsymbol{p} \in \widetilde{T}$, we remove the at most $k(n-1)-1$ other sets from $\widetilde{\mathcal{T}}$ that contain $\boldsymbol{p}$. We continue with the next set in $\widetilde{\mathcal{T}}$ that has not yet been removed until $\widetilde{\mathcal{T}} = \emptyset$. Let $\mathcal{T}^\star \subseteq \widetilde{\mathcal{T}}$ be the family of sets that we obtain by this process. Clearly, $\mathcal{T}^\star$ is a Tverberg partition and because $\mathcal{T}^\star \subseteq \widetilde{\mathcal{T}}$, we have $\boldsymbol{p}^\star \in \bigcap_{\widetilde{T}\in\mathcal{T}^\star} \operatorname{conv}(\widetilde{T})$. Moreover, for each set $\widetilde{T}_i \in \mathcal{T}^\star$, we remove at most $(d+1)(k(n-1)-1)$ other sets from $\widetilde{\mathcal{T}}$. Thus, the size of the Tverberg partition $\mathcal{T}^\star$ is at least

$$|\mathcal{T}^\star| \geq \left\lceil \frac{m}{(d+1)(k(n-1)-1)+1} \right\rceil \geq \left\lceil \frac{n}{(d+1)^2(k(n-1)-1)+d+1} \right\rceil.$$

Constructing the ColorfulCarathéodory instance takes $O(mn^2)$ time. Using $\mathcal{A}$, we need $T(n-1)$ time to compute a $k(n-1)$-colorful choice $\widetilde{C}$. Pruning every set of $\widetilde{\mathcal{T}}$ with Theorem 1.1 to at most $d+1$ points needs $O(m(d^3n+n^2)) = O((d^2+m)n^2)$ time. Finally, constructing $\mathcal{T}^\star$ out of $\widetilde{\mathcal{T}}$ takes $O(n^2)$ time with the naive algorithm. This results in the claimed running time of $O((d^2+m)n^2+T(n-1))$. ■

The next corollary is a direct consequence of Corollary 7.12 and Lemma 1.11.

Corollary 7.13. *Let $P \subset \mathbb{R}^d$ be a set of size n and let $\mathcal{A}$ be as above. Furthermore, let $\widetilde{m}$ be as in Corollary 7.12. Then, we can compute a point with simplicial depth at least*

$$\left\lceil \frac{\widetilde{m}^{d+1}}{(d+1)^{d+1}} \right\rceil = \Omega\left(\frac{n^{d+1}}{d^{3d+3}\,(k(n-1))^{d+1}}\right)$$

with respect to P.

Proof. Using the algorithm from Corollary 7.12, we can compute a Tverberg m-partition $\mathcal{T}$ for P together with a point $\boldsymbol{p}^\star \in \bigcap_{T\in\mathcal{T}} \operatorname{conv}(T)$ in $O((d^2+m)n^2+T(n-1))$ time. Then, the point $\boldsymbol{p}^\star$ has simplicial depth at least

$$\left\lceil \frac{\widetilde{m}^{d+1}}{(d+1)^{d+1}} \right\rceil = \Omega\left(\frac{n^{d+1}}{d^{3d+3}\,(k(n-1))^{d+1}}\right)$$

with respect to P by Lemma 1.11. ■

Furthermore, we can use $\mathcal{A}$ to approximate ColorfulKirchberger. In Theorem 1.10, we are only allowed to take one point from each Tverberg partition. Here however, we allow to take multiple points from each Tverberg partition.

Corollary 7.14. *Let $\mathcal{A}$ be as above and let $C_1, \dots, C_n \subset \mathbb{R}^d$ be $n = (m-1)(d+1)+1$ pairwise disjoint color classes that are each of size n. Furthermore, for $i \in [n]$, let $\mathcal{T}_i = \{T_{i,1}, \dots, T_{i,m}\}$*

denote a Tverberg m-partition for C_i. Then, given for each Tverberg partition $\mathcal{T}_i$, $i \in [n]$, a point $\boldsymbol{p}_i \in \bigcap_{j=1}^{m} \operatorname{conv}(T_{i,j})$, and for all $i \in [n]$ and $j \in [m]$, the coefficients of the convex combination of $\boldsymbol{p}_i$ with the points in $T_{i,j}$, we can compute in $O\big(n^3 + T(n-1)\big)$ time a $k(n-1)$-colorful choice $C \subseteq \bigcup_{i=1}^{n} C_i$ such that

$$\mathcal{T}_C = \left\{ C \cap \left(\bigcup_{i=1}^{n} T_{i,j} \right) \middle| \, j \in [m] \right\}$$

is a Tverberg m-partition for C.

Proof. In the proof of Theorem 1.10 on page 12, we lift the points $\bigcup_{i=1}^{n} C_i$ to $\mathbb{R}^{n-1}$ such that the set of points $\widehat{C}_i$ that corresponds to the color class C_i still embraces the origin, where $i \in [n]$. Moreover, if $\widehat{C}' \subseteq \bigcup_{i=1}^{n} \widehat{C}_i$ is a $\mathbf{0}$-embracing colorful choice of the lifted points, then there is a corresponding colorful choice C' with respect to $C_1, \dots, C_n$ such that

$$\mathcal{T}_{C'} = \left\{ C' \cap \left(\bigcup_{i=1}^{n} T_{i,j} \right) \middle| \, j \in [m] \right\}$$

is a Tverberg m-partition for C'. Similarly, a $\mathbf{0}$-embracing $k(n-1)$-colorful choice $\widehat{C}$ of the lifted color classes corresponds to a $k(n-1)$-colorful choice C with respect to $C_1, \dots, C_n$ such that

$$\mathcal{T}_C = \left\{ C \cap \left(\bigcup_{i=1}^{n} T_{i,j} \right) \middle| \, j \in [m] \right\}$$

is a Tverberg m-partition for C.

Computing the function $\boldsymbol{p} \otimes \boldsymbol{q}$, where $\boldsymbol{p} \in \mathbb{R}^{d+1}$ and $\boldsymbol{q} \in \mathbb{R}^{m-1}$, needs $O(dm) = O(n)$ time and hence lifting the point sets $C_1, \dots, C_n \subset \mathbb{R}^d$ to $\mathbb{R}^{n-1}$ with Lemma 1.9 needs $O\big(n^3\big)$ time in total. Since we know for each Tverberg partition $\mathcal{T}_i$, $i \in [n]$, a point $\boldsymbol{p}_i \in \bigcap_{j=1}^{m} \operatorname{conv}\big(T_{i,j}\big)$ together with the coefficients of the convex combination of $\boldsymbol{p}_i$ with the points in $T_{i,j}$ for $j \in [m]$, we can compute in $O(n)$ time the coefficients of the convex combination of the origin with the points in $\widehat{C}_i$ as described in the proof of Lemma 1.9. Then, $\mathcal{A}$ can be applied to compute a $\mathbf{0}$-embracing $k(n-1)$-colorful choice $\widehat{C}$ with respect to the lifted point sets in $T(n-1)$ time. Finally, constructing C and $\mathcal{T}_C$ out of $\widehat{C}$ needs $O(n)$ time. Hence, the total time needed is $O\big(n^3 + T(n-1)\big)$. ■

Now, given $d+1$ color classes $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ that embrace the origin, we can compute with Algorithm 7.2 an $\lceil \varepsilon d \rceil$-colorful choice that embraces the origin in polynomial time. Combining this with Corollary 7.12, we obtain an algorithm that computes Tverberg partitions of size $O(1)$ in polynomial time, a trivial result. Since Corollary 7.13 reduces to Corollary 7.12, we also do not obtain a nontrivial approximation algorithm for SIMPLICIALCENTER. However, combining Algorithm 7.2 with Corollary 7.14, we do obtain a nontrivial approximation algorithm for COLORFULKIRCHBERGER: given $n = (m-1)(d+1)+1$ color classes $C_1, \dots, C_n$, each of size n, and for each color class a Tverberg m-partition $\mathcal{T}_i = \{T_{i,1}, \dots, T_{i,m}\}$ together with a point $\boldsymbol{p}_i \in \bigcap_{j=1}^{m} \operatorname{conv}\big(T_{i,j}\big)$ and the coefficients of the convex combination of $\boldsymbol{p}_i$ with the points in

$T_{i,j}$ for all $j \in [m]$, we can compute in $n^{O(\varepsilon^{-1} \ln \varepsilon^{-1})}$ time an $\lceil \varepsilon n \rceil$-colorful choice C such that

$$\mathcal{T}_C = \left\{ C \cap \left(\bigcup_{i=1}^{n} T_{i,j} \right) \middle| \, j \in [m] \right\}$$

is a Tverberg m-partition for C, where $\varepsilon > 0$ is arbitrary but fixed.

8 Few Color Classes

In the last chapter, we relaxed the problem of finding a $\mathbf{0}$-embracing colorful choice by taking multiple points from each color class. In this chapter, we consider a further natural relaxation of ColorfulCarathéodory that is closely related to k-colorful choices. Let $C_1, \ldots, C_m \in \mathbb{Q}^d$ be m color classes and let $C \subseteq \bigcup_{i=1}^m C_i$ be a set. We call C an $(k_1, \ldots, k_m)$*-colorful choice* with respect to $C_1, \ldots, C_m$ if there are m subsets $C_1' \subseteq C_1, \ldots, C_m' \subseteq C_m$ with $|C_i'| \le k_i$ for $i \in [m]$ such that $C = \bigcup_{i=1}^m C_i'$. Now, given m color classes $C_1, \ldots, C_m \subset \mathbb{Q}^d$ that each ray-embrace a point $\boldsymbol{b} \in \mathbb{Q}^d$, $\boldsymbol{b} \neq \mathbf{0}$, and given m numbers $k_1, \ldots, k_m \in \mathbb{N}$ that sum to d, we want to find a $(k_1, \ldots, k_m)$-colorful choice that ray-embraces $\boldsymbol{b}$. It is a straightforward consequence of the colorful Carathéodory theorem that a $(k_1, \ldots, k_m)$-colorful choice that ray-embraces $\boldsymbol{b}$ always exists.

Corollary 8.1 (of Theorem 1.2). *Let $C_1, \ldots, C_m \subset \mathbb{R}^d$ be m finite sets that ray-embrace a point $\boldsymbol{b} \in \mathbb{R}^d$, $\boldsymbol{b} \neq \mathbf{0}$. Furthermore, let $k_1, \ldots, k_m \in \mathbb{N}$ be m numbers that sum to d. Then, there exists an $(k_1, \ldots, k_m)$-colorful choice that ray-embraces $\boldsymbol{b}$.*

Proof. Let $C_1, \ldots, C_m \in \mathbb{R}^d$ be m color classes that ray-embrace $\boldsymbol{b} \in \mathbb{R}^d$ and let $k_1, \ldots, k_m \in \mathbb{N}$ be m numbers that sum to d. We set $C_{i,1}, \ldots, C_{i,k_i}$ to C_i for $i \in [m]$. Then, we have d color classes $C_{1,1}, \ldots, C_{1,k_1}, \ldots, C_{m,1}, \ldots, C_{m,k_m}$ that each ray-embrace $\boldsymbol{b}$. Thus by Theorem 1.2, there are d points $\boldsymbol{c}_{1,1} \in C_{1,1}, \ldots, \boldsymbol{c}_{1,k_1} \in C_{1,k_1}, \ldots, \boldsymbol{c}_{m,1} \in C_{m,1}, \ldots, \boldsymbol{c}_{m,k_m} \in C_{m,k_m}$ that ray-embrace $\boldsymbol{b}$. Set $C_i' = \{\boldsymbol{c}_{i,j} \mid j \in [k_i]\}$ for $i \in [m]$. Then, the set $\bigcup_{i=1}^m C_i'$ is a $(k_1, \ldots, k_m)$-colorful choice that ray-embraces $\boldsymbol{b}$. ■

Clearly, for $m = d$, the problem coincides with ColorfulCarathéodory. Surprisingly, even if $m = 2$ and we are allowed to take half of each of the two color classes, it is already nontrivial to find a solution. Using our techniques from Chapter 4, we present a weakly polynomial-time algorithm on a Word-Ram for this case. As described in Section 4.2, 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.

In the following, we use the same notation as in Chapter 4 (see Table 4.1 on page 46 for an overview). Let $C_1, C_2 \subset \mathbb{Q}^d$ be two color classes, each of size d, let $\boldsymbol{b} \in \mathbb{Q}^d$, $\boldsymbol{b} \neq \mathbf{0}$, 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 Chapter 3 and Chapter 4, we introduce $d-2$ "dummy" color classes $C_3, \ldots, C_d$ that trivially ray-embrace $\boldsymbol{b}$ by setting $C_3 = \cdots = C_d = \{\boldsymbol{b}\}$. Let $(C_1', \ldots, C_d', \boldsymbol{b}')$ be the instance of ColorfulCarathéodory in general position that we obtain by applying Lemma 3.6 to $(C_1, \ldots, C_d, \boldsymbol{b})$. Then, let $\mathcal{P}^{\mathrm{CC}} \subset \mathbb{Q}^{d^2}$ denote the polyhedron that is defined by the linear system L^{CC} (see (4.1) on page 36) for the instance $(C_1', \ldots, C_d', \boldsymbol{b}')$. 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 Corollary 4.9, 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 8.2. *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 Corollary 4.8, 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 4.12 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 Corollary 4.8, 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 4.11 implies the statement. ■

Using Lemma 8.2, 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 8.3. *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 8.2, 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 Corollary 4.8, $f^\star$ 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 8.3 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 8.3, 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 Corollary 4.8, $\boldsymbol{\mu}' = \Phi_\Delta(f') \cap g$ 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 8.2, 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 8.3 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 8.4. *Let L be the length of the binary encoding of the* ColorfulCarathéodory *instance $(C_1', \dots, C_d', \boldsymbol{b}')$ 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 Corollary 4.8, 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 2.3 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 $(C_1, \dots, C_d, \boldsymbol{b})$, 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 3.5 then guarantees that C ray-embraces $\boldsymbol{b}$.

Theorem 8.5. *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 $\boldsymbol{b}$ in weakly polynomial time on a* Word-Ram.

For Sperner's lemma, it is well-known that a fully-labeled simplex can be found by binary search if there are only two colors. And this is essentially 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.

9 Exact Algorithms for COLORFULCARATHÉODORY

In contrast to the past chapters, we now focus on the problem of computing an exact solution for the convex version of COLORFULCARATHÉODORY. Let $C_1, \dots, C_{d+1} \subset \mathbb{Q}^d$ be $d+1$ sets that embrace the origin and assume all are of size at most $d+1$. The naive algorithm checks for all $O\left(d^{d+1}\right)$ possible colorful choices whether they embrace the origin. This can be further improved by using the following result by Bárány.

Theorem 9.1 ([9, Theorem 2.3])**.** *Let $C_1, \dots, C_d \subset \mathbb{R}^d$ be d sets that all embrace the origin and let $\boldsymbol{c} \in \mathbb{R}^d$ be a point. Then, there exist d points $\boldsymbol{c}_1 \in C_1, \dots, \boldsymbol{c}_d \in C_d$ such that the set $\{\boldsymbol{c}, \boldsymbol{c}_1, \dots, \boldsymbol{c}_d\}$ embraces the origin.* ■

In particular, Theorem 9.1 implies that every point $\boldsymbol{c} \in \bigcup_{i=1}^{d+1} C_i$ participates in some $\mathbf{0}$-embracing colorful choice and hence we can fix a point from one color class and check only all $O\left(d^d\right)$ possibilities of extending it to a colorful choice.

We now consider two related settings that allow for further improvement. We begin with the simple case in which each color class consists of only two points. Then basic linear algebra suffices to compute a $\mathbf{0}$-embracing colorful choice in polynomial-time. In Section 9.2, we show that many color classes help. Using an approach similar to the algorithm by Miller and Sheehy for approximating Tverberg partitions [57], we present a quasi-polynomial time algorithm that computes a $\mathbf{0}$-embracing colorful choice when given $\Theta\left(d^2 \log d\right)$ color classes instead of only $d+1$.

9.1. A Simple Special Case

In the following, we assume that $|C_1| = \cdots = |C_{d+1}| = 2$ and let $\boldsymbol{c}_{i,1}, \boldsymbol{c}_{i,2}$ denote the two points in C_i for $i \in [d+1]$. Clearly, for all $i \in [d+1]$, the point $-\boldsymbol{c}_{i,1}$ must be contained in the positive span of $\boldsymbol{c}_{i,2}$. Furthermore, we assume without loss of generality that all points are different from the origin, as otherwise computing a $\mathbf{0}$-embracing colorful choice is trivial. Then, the set $\left\{\boldsymbol{c}_{i,1} \,\middle|\, i \in [d+1]\right\}$ is linearly dependent and hence there exist coefficients $\phi_1, \dots, \phi_{d+1} \in \mathbb{R}$, not all 0, such that $\mathbf{0} = \sum_{i=1}^{d+1} \phi_i \boldsymbol{c}_{i,1}$. Now, since $-\boldsymbol{c}_{i,1} \in \operatorname{pos}\left(\boldsymbol{c}_{i,2}\right)$ for all $i \in [d+1]$, the set $C = \left\{\boldsymbol{c}_{i,1} \,\middle|\, i \in [d+1], \phi_i \geq 0\right\} \cup \left\{\boldsymbol{c}_{i,2} \,\middle|\, i \in [d+1], \phi_i < 0\right\}$ embraces the origin and it is a colorful choice. Since the computation of the coefficients of the linear dependency can be carried out in $O\left(d^3\right)$ time with Gaussian elimination, finding C takes $O\left(d^3\right)$ time in total. The following theorem is now immediate.

Theorem 9.2. *Let $C_1, \dots, C_{d+1} \subset \mathbb{R}^d$ be $d+1$ pairs of points that all embrace the origin. Then, a $\mathbf{0}$-embracing colorful choice can be computed in $O\left(d^3\right)$ time.*

9.2. Many Colors

In the following, we assume that we are given $\Theta\left(d^2 \log d\right)$ instead of only $d+1$ color classes that all embrace the origin. The algorithm repeatedly combines k-colorful choices to one $\mathbf{0}$-embracing $\lceil k/2 \rceil$-colorful choice until a $\mathbf{0}$-embracing 1-colorful choice is obtained. This approach is similar to the Miller-Sheehy approximation algorithm for Tverberg partitions [57] and leads to an algorithm with total running time $d^{O(\log d)}$.

Lemma 9.3. *Let $C'_1, \ldots, C'_{d+1} \subset \mathbb{R}^d$ be $\mathbf{0}$-embracing k-colorful choices of size $O(d)$ such that each color appears in a unique k-colorful choice. Then, given the coefficients of the convex combination of the origin for each set C'_i, $i \in [d+1]$, a $\mathbf{0}$-embracing $\lceil k/2 \rceil$-colorful choice C' can be computed in $O\left(d^5\right)$ time.*

Proof. First, we prune each k-colorful choice C'_i, $i \in [d+1]$, with Lemma 7.4 and then partition it into two sets $C'_{i,1}, C'_{i,2}$ by distributing the points from each color equally among both sets. Then, we apply the algorithm from Lemma 7.6 to obtain two representative points $\boldsymbol{r}_{i,1}, \boldsymbol{r}_{i,2}$ and set $R_i = \{\boldsymbol{r}_{i,1}, \boldsymbol{r}_{i,2}\}$. Since the sets $R_1, \ldots, R_{d+1}$ each embrace the origin and consist of only two points, we can compute a 1-colorful choice R with respect to $R_1, \ldots, R_{d+1}$ with the algorithm from Theorem 9.2. Now, consider the set $C' = \left\{C'_{i,j} \,\middle|\, \boldsymbol{r}_{i,j} \in R\right\}$. Since R is $\mathbf{0}$-embracing, so is C'. Moreover, because a color j appears only in one of the k-colorful choices, say C'_i, and since each set of the partition $C'_{i,1}, C'_{i,2}$ contains at most $\lceil k/2 \rceil$ points with color j, the set C' is a $\lceil k/2 \rceil$-colorful choice.

Pruning each k-colorful choice with Lemma 7.4 and then computing the two representative points per partition takes $O\left(d^5\right)$ time in total. This dominates the time needed for the computation of R and thus, we can compute C' in $O\left(d^5\right)$ time. ∎

Note that Lemma 9.3 actually implies a second algorithm to compute $\lceil (d+1)/2 \rceil$-colorful choices that embrace the origin: let $C_1, \ldots, C_{d+1} \subset \mathbb{R}^d$ be $\mathbf{0}$-embracing color classes and assume the sets have size $d+1$. Set $C'_i = C_i$ in Lemma 9.3 for $i \in [d+1]$. Then, C'_i is trivially a $(d+1)$-colorful choice and hence the set C' is a $\lceil (d+1)/2 \rceil$-colorful choice.

Now, we apply Lemma 9.3 repeatedly until we obtain a 1-colorful choice as follows. Let $C_1, \ldots, C_n \subset \mathbb{Q}^d$ be $n = \Theta\left(d^2 \log d\right)$ color classes such that C_i is $\mathbf{0}$-embracing and has size $O(d)$ for $i \in [n]$. We create an array A of size $m = \Theta(\log d)$ that initially contains all n color classes in $A[0]$. Set $c_0 = d+1$ and for $i \in [k]$, set $c_i = \lceil c_{i-1}/2 \rceil$. Throughout the algorithm, we maintain the invariant that the ith cell contains only $\mathbf{0}$-embracing c_i-colorful choices and that each color appears in at most one set in all of A. Since $c_0 = d+1$, the invariant holds in the beginning. We repeatedly improve k-colorful choices with Lemma 9.3 as follows: let i be the maximum index of a cell in A that contains at least $d+1$ sets $C'_1, \ldots, C'_{d+1}$ and remove them from $A[i]$. By our invariant, these sets are $\mathbf{0}$-embracing c_i-colorful choices. Applying Lemma 9.3, we can combine $C'_1, \ldots, C'_{d+1}$ to one c_{i+1}-colorful choice C' that embraces the origin. We prune it with Lemma 7.4 and check whether it is a 1-colorful choice. If so, we have found a solution. Otherwise, we add it to $A[i+1]$. Furthermore, we check for colors that appeared in the removed sets $C'_1, \ldots, C'_{d+1}$ but not in C' and add the corresponding color classes back to $A[0]$. The invariant is maintained since these colors only appeared in the removed sets. See Algorithm 9.1 for a detailed description of the algorithm.

Algorithm 9.1: Exact algorithm for many color classes.

Input: color classes $C_1, \ldots, C_n \subset \mathbb{R}^d$ and for each set C_i, the coefficients of the convex combination of $\mathbf{0}$, where $n = \Theta\left(d^2 \log d\right)$

$A \leftarrow$ Array of size $m = \Theta\left(\log d\right)$;
$A[0] \leftarrow \{C_1, \ldots, C_n\}$;
while *no* $\mathbf{0}$*-embracing colorful choice was found* **do**
    $i \leftarrow$ maximum index with $|A[i]| \geq d+1$;
    Remove $d+1$ sets $C'_1, \ldots, C'_{d+1}$ from $A[i]$;
    $C' \leftarrow$ combine $C'_1, \ldots, C'_{d+1}$ with Lemma 9.3;
    Prune C' with Lemma 7.4;
    if C' *is a colorful choice* **then**
        return C';
    Add C' to $A[i+1]$;
    Add all color classes C_i with $C_i \cap \left(\bigcup_{i=1}^{d+1} C'_i\right) \neq \emptyset$ and $C_i \cap C' = \emptyset$ to $A[0]$;

We conclude this chapter by proving correctness of Algorithm 9.1 and analyzing its running time.

Theorem 9.4. *Let $C_1, \ldots, C_n \subset \mathbb{R}^d$ be $n = \Theta\left(d^2 \log d\right)$ sets such that C_i embraces the origin and $|C_i| = O(d)$, where $i \in [n]$. Then, given the coefficients of the convex combination of the origin for each set C_i, $i \in [n]$, Algorithm 9.1 computes a $\mathbf{0}$-embracing colorful choice in $d^{O(\log d)}$ time.*

Proof. We have already argued that the ith cell of the array A contains only $\mathbf{0}$-embracing c_i-colorful choices. First, we observe that progress is always possible, i.e., that it is always possible to find a cell of A that contains at least $d+1$ sets: the array has $m = \Theta\left(\log d\right)$ levels and within each set in A, at most d colors appear. Thus, for $d^2 m + 1 = \Theta\left(d^2 \log d\right)$ colors, the pigeonhole principle guarantees a cell with at least $d+1$ sets.

Set now $m = \lceil \log(d+1) \rceil + 1$. We claim that a combination of $d+1$ sets in $A[m]$ results in a $\mathbf{0}$-embracing colorful choice. Now, since $c_i \leq \frac{d+1}{2^i} + 2$, the sets in $A[m-1]$ are $\mathbf{0}$-embracing 3-colorful choices, the sets in $A[m]$ are 2-colorful choices and the combination of $d+1$ sets in $A[m]$ gives a 1-colorful choice, as claimed.

To compute a set in level i, we have to compute $d+1$ sets in level $i-1$. Since one application of Lemma 9.3 takes $O\left(d^5\right)$ time, we need $d^{O(i)}$ time in total to compute a set in level i and hence the total running time of the algorithm is bounded by $d^{O(\log d)}$. ■

10 Simplicial Depth

The simplicial depth $\sigma_P(q)$ of a point $\boldsymbol{q} \in \mathbb{R}^d$ with respect to a point set $P \subset \mathbb{R}^d$ is the number of distinct d-simplices with vertices in P that embrace $\boldsymbol{q}$. As presented in Chapter 1, the colorful Carathéodory theorem implies the existence of a point with simplicial depth $\Theta\left(n^{d+1}\right)$ if d is fixed. In this chapter, we consider the problem of computing the simplicial depth of a given point.

For two dimensions, there are several algorithms that compute the simplicial depth $\sigma_P(q)$ in $O\left(n \log n\right)$ time [33, 41, 68], where n is the size of P. Cheng and Ouyang presented an algorithm for three dimensions that needs $O\left(n^2\right)$ time and generalized it to four dimensions at the cost of an increased running time of $O\left(n^4\right)$ [22]. For higher but fixed dimensions, there is no published improvement on the naive $\Theta\left(n^{d+1}\right)$ algorithm that tests all possible d-simplices. We show in Section 10.3 how the running time of the naive algorithm can be improved to $O\left(n^d \log n\right)$ time by using a simplex range searching data structure. Recently, although yet unpublished, this was further improved by Pilz et al. [62] who, based on a novel dimension reduction argument, presented an algorithm that computes the simplicial depth in $O\left(n^d\right)$ time.

In Section 10.2, we present two algorithms that each compute a $(1+\varepsilon)$-approximation of the simplicial depth, however with different worst-case scenarios. A combination of these strategies gives an algorithm that returns a $(1+\varepsilon)$-approximation of the simplicial depth with high probability in $\widetilde{O}\left(n^{d/2+1}\right)$ time. Finally, we show in Section 10.4 that computing the simplicial depth becomes #P-complete and W[1]-hard with respect to the parameter d if the dimension is part of the input. This directly implies W[1]-hardness of computing the *colorful simplicial depth*: here, we are given m color classes $C_1, \dots, C_m \subset \mathbb{R}^d$ and a point $\boldsymbol{q} \in \mathbb{R}^d$ and we want to count the number $\sigma^{\text{col}}_{C_1,\dots,C_m}(q)$ of $(d+1)$-sets that contain at most one point from each color and that embrace $\boldsymbol{q}$. By assigning each point in P its own color, we can trivially reduce the problem of computing the simplicial depth to the problem of computing the colorful simplicial depth. For the colorful simplicial depth, #P-completeness is already known [11, Theorem 5.4]. We begin with a discussion on the Gale transform, a notion of duality which we will use extensively in the remainder of this chapter, and the complexity classes #P and W[1].

10.1. Preliminaries

10.1.1. The Gale Transform

The Gale transform (see e.g. [48, Section 5.6]) maps a sequence of n points P in $\mathbb{R}^d$ to a sequence GT(P) of n points in $\mathbb{R}^{n-d-1}$ such that there is a bijection between the facets of the polytope conv(P) and the $\mathbf{0}$-embracing simplices with vertices in GT(P). We assume that conv(P) is simplicial and that $\dim P = d$. Then, the construction of the Gale transform for $P = \{\boldsymbol{p}_1, \dots, \boldsymbol{p}_n\} \subset \mathbb{R}^d$ is defined as follows.

$$A = \left(\begin{array}{|ccc|} \hline & & \\ & P & \\ & & \\ \hline 1 & \dots & 1 \end{array}\right) = \left(\begin{array}{|c|c|c|c|} \hline \hat{\boldsymbol{p}}_1 & \hat{\boldsymbol{p}}_2 & \cdots & \hat{\boldsymbol{p}}_n \\ \hline \end{array}\right) = \left(\begin{array}{|c|} \hline \boldsymbol{r}_1 \\ \hline \boldsymbol{r}_2 \\ \hline \vdots \\ \hline \boldsymbol{r}_{d+1} \\ \hline \end{array}\right) \in \mathbb{R}^{(d+1)\times n}$$

$$G = \left(\begin{array}{|c|c|c|c|} \hline \boldsymbol{g}_1 & \boldsymbol{g}_2 & \cdots & \boldsymbol{g}_{n-d-1} \\ \hline \end{array}\right) = \left(\begin{array}{|c|} \hline \overline{\boldsymbol{p}}_1 \\ \hline \overline{\boldsymbol{p}}_2 \\ \hline \vdots \\ \hline \overline{\boldsymbol{p}}_n \\ \hline \end{array}\right) \in \mathbb{R}^{n\times(n-d-1)}$$

Figure 10.1.: Overview of the notation for the Gale transform.

1. We lift P to a sequence $\widehat{P} \subset \mathbb{R}^{d+1}$ by appending a 1-coordinate to each point:
$$\widehat{P} = \left\{ \hat{\boldsymbol{p}}_i = \begin{pmatrix} \boldsymbol{p}_i \\ 1 \end{pmatrix} \,\middle|\, \boldsymbol{p}_i \in P \right\} \subset \mathbb{R}^{d+1}.$$
2. Let $A = (\hat{\boldsymbol{p}}_1 \dots \hat{\boldsymbol{p}}_n) \in \mathbb{R}^{(d+1)\times n}$ denote the matrix whose ith column vector is $\hat{\boldsymbol{p}}_i$ and let $\boldsymbol{r}_i \in \mathbb{R}^n$ denote the ith row vector of A. Then, we take a basis $\boldsymbol{g}_1, \dots, \boldsymbol{g}_{n-d-1} \in \mathbb{R}^n$ of the orthogonal complement of span$(\boldsymbol{r}_1, \dots, \boldsymbol{r}_{d+1})$. Please see Figure 10.1 for an overview of the notation.
3. Let $G = (\boldsymbol{g}_1 \dots \boldsymbol{g}_{n-d-1}) \in \mathbb{R}^{n\times(n-d-1)}$ denote the matrix with $\boldsymbol{g}_i$ as ith column vector and let $\overline{\boldsymbol{p}}_i \in \mathbb{R}^{n-d-1}$ denote the ith row vector of G. Then, we set $\mathrm{GT}(P) = \{\overline{\boldsymbol{p}}_1, \dots, \overline{\boldsymbol{p}}_n\}$.

Lemma 10.1 (Gale transform)**.** *Let $P = \{\boldsymbol{p}_1, \dots, \boldsymbol{p}_n\} \subset \mathbb{R}^d$ be a sequence of points such that* conv(P) *is simplicial and* $\dim P = d$. *Furthermore, let* $\mathrm{GT}(P) = \{\overline{\boldsymbol{p}}_1, \dots, \overline{\boldsymbol{p}}_n\}$ *denote the Gale transform of P. Then, a set $P' \subseteq P$, $|P'| = d$, defines a facet of* conv(P) *if and only if the simplex*

$$\mathrm{conv}\left\{\overline{\boldsymbol{p}}_i \in \mathrm{GT}(P) \,\middle|\, \boldsymbol{p}_i \in P \setminus P'\right\} \subset \mathbb{R}^{n-d-1}$$

embraces the origin.

Proof. Let P' be a d-subset of P that defines a facet of conv(P) and let $\boldsymbol{p}_{k_1}, \dots, \boldsymbol{p}_{k_d}$ denote the points in P', where $k_1, \dots, k_d \in [n]$. Furthermore, let

$$\widehat{P'} = \left\{ \hat{\boldsymbol{p}}_{k_i} = \begin{pmatrix} \boldsymbol{p}_{k_i} \\ 1 \end{pmatrix} \,\middle|\, i \in [d] \right\} \subset \mathbb{R}^{d+1}$$

denote the points from P' lifted to $\mathbb{R}^{d+1}$ by appending a 1-coordinate. Then, all points of P are on one side of the hyperplane $h = \operatorname{aff} P'$ and hence all points from $\widehat{P}$ are one side of the hyperplane $\hat{h} = \operatorname{span} \widehat{P}'$. Let $\hat{\boldsymbol{v}} \in \mathbb{R}^{d+1}$, $\hat{\boldsymbol{v}} \neq \mathbf{0}$, be the normal vector to $\hat{h}$ and assume without loss of generality that

$$\hat{\boldsymbol{v}}^T \hat{\boldsymbol{p}}_i \geq 0 \text{ for } i \in [n], \tag{10.1}$$

and that $\hat{\boldsymbol{v}}^T \hat{\boldsymbol{p}}_i = 0$ for $\boldsymbol{p}_i \in P'$. Consider now the coordinates of $\hat{\boldsymbol{v}}$ as coefficients of a linear combination of the row vectors $\boldsymbol{r}_1, \ldots, \boldsymbol{r}_{d+1}$ and let $\boldsymbol{r} \in \mathbb{R}^n$ denote its result, i.e., set

$$\boldsymbol{r} = \sum_{i=1}^{d+1} (\hat{\boldsymbol{v}})_i \boldsymbol{r}_i = \begin{pmatrix} \hat{\boldsymbol{v}}^T \hat{\boldsymbol{p}}_1 & \ldots & \hat{\boldsymbol{v}}^T \hat{\boldsymbol{p}}_1 \end{pmatrix}^T \in \mathbb{R}^n.$$

Because of (10.1), all coordinates of $\boldsymbol{r}$ are nonnegative and $(\boldsymbol{r})_i = 0$ for $i \in \{k_1, \ldots, k_d\}$. Furthermore, since $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_{n-d-1}$ is a basis of the orthogonal complement of $\operatorname{span}\{\boldsymbol{r}_1, \ldots, \boldsymbol{r}_{d+1}\}$, we have

$$0 = \boldsymbol{r}^T \boldsymbol{g}_i = \left(\sum_{j=1}^{n} (\boldsymbol{r})_j \left(\overline{\boldsymbol{p}}_j \right)_i \right) \text{ for } i \in [n-d-1].$$

Thus, thinking of the coordinates of $\boldsymbol{r}$ as the coefficients of a positive combination of the points in $\mathrm{GT}(P)$, we obtain

$$\sum_{i=1}^{n} (\boldsymbol{r})_i \overline{\boldsymbol{p}}_i = \mathbf{0},$$

and hence $\mathrm{GT}(P)$ embraces the origin. Moreover, since $(\boldsymbol{r})_i = 0$ for $i \in \{k_1, \ldots, k_d\}$, the simplex $\operatorname{conv}\left\{\overline{\boldsymbol{p}}_i \in \mathrm{GT}(P) \,\middle|\, \boldsymbol{p}_i \notin P'\right\}$ embraces the origin. Since the above arguments are equivalences, the statement follows. ■

Since the computation of $\mathrm{GT}(P)$ involves only finding a basis for the orthogonal complement of a linear subspace, it can be carried out in $O\left(d^2 n\right)$ time [7, Lemma 1]. Furthermore, if the origin is the barycenter of P then there always exists a sequence of points $\mathrm{GT}^{-1}(P) \subset \mathbb{R}^{n-d-1}$ whose Gale transform results in P. We can think of these points as inverse Gale transform of P.

Lemma 10.2 (Inverse Gale transform)**.** *Let $P \subset \mathbb{R}^d$ be a point sequence of size n such that the origin is the barycenter of P and* $\dim P = d$*. Then, there exists a point sequence* $\mathrm{GT}^{-1}(P) \subset \mathbb{R}^{n-d-1}$ *such that* $P = \mathrm{GT}\left(\mathrm{GT}^{-1}(P)\right)$.

Proof. We can easily construct a point sequence $\mathrm{GT}^{-1}(P) \subset \mathbb{R}^{n-d-1}$ whose Gale transform results in P by inverting the steps of the Gale transform as follows. Let $\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n$ denote the points in P and let

$$G = \begin{pmatrix} \boldsymbol{p}_1^T \\ \vdots \\ \boldsymbol{p}_n^T \end{pmatrix} \in \mathbb{R}^{n \times d}$$

be the matrix with the points from P as row vectors. Let $\boldsymbol{g}_i \in \mathbb{R}^n$, $i \in [d]$, denote the ith column vector of G. Since the origin is the barycenter of P, we have

$$\mathbf{1}^T \boldsymbol{g}_i = \sum_{j=1}^{n} \left(\boldsymbol{p}_j \right)_i = 0 \text{ for } i \in [n],$$

and hence the vector $\mathbf{1}^T \in \mathbb{R}^n$ is orthogonal to $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_d$. We set $\boldsymbol{r}_{n-d} = \mathbf{1}^T \in \mathbb{R}^n$ and complete it to a basis $\boldsymbol{r}_1, \ldots, \boldsymbol{r}_{n-d}$ of the orthogonal complement of span$\{\boldsymbol{g}_1, \ldots, \boldsymbol{g}_d\}$. Now, let

$$A = \begin{pmatrix} \boldsymbol{r}_1^T \\ \vdots \\ \boldsymbol{r}_{n-d}^T \end{pmatrix} \in \mathbb{R}^{(n-d)\times n}$$

denote the matrix with $\boldsymbol{r}_1, \ldots, \boldsymbol{r}_{n-d}$ as row vectors. Let $\hat{\boldsymbol{p}}_i = \begin{pmatrix} \boldsymbol{p}_i & 1 \end{pmatrix}^T \in \mathbb{R}^{n-d}$, $i \in [n]$, denote the ith column vector of A and set $\mathrm{GT}^{-1}(P) = \{\boldsymbol{p}_i \mid i \in [n]\}$. Then, $P = \mathrm{GT}\left(\mathrm{GT}^{-1}(P)\right)$, as claimed. ■

10.1.2. The Complexity Class #P

Let $\mathcal{L} \subseteq \{0,1\}^\star$ be a language. Then, $\mathcal{L}$ is in NP if and only if there exists a polynomial p and a polynomial-time Turing machine M, the *verifier*, such that for all words $w \in \mathcal{L}$ there exists a word $c \in \{0,1\}^{p(|w|)}$ such that M returns on input (w,c) YES. We call c a *certificate* for w. Furthermore, for all words $w \notin \mathcal{L}$ and for all words $c \in \{0,1\}^{p(|w|)}$, the Turing machine M returns on input (w,c) NO. Deciding whether a word $w \in \{0,1\}^\star$ is in $\mathcal{L}$ is equivalent to deciding whether there is *at least one* certificate for w. Now, the complexity class #P [79] captures the complexity of counting *how many* certificates there are for w. More formally, the complexity class #P contains functions $f : \{0,1\}^\star \mapsto \mathbb{N}_0$ for which there exists a polynomial p_f and a polynomial-time Turing machine M_f such that for all words $w \in \{0,1\}^\star$, we have

$$f(w) = \left|\left\{c \,\middle|\, c \in \{0,1\}^{p_f(|w|)}, M_f \text{ returns on input } (w,c) \text{ YES}\right\}\right|.$$

To define #P-completeness, we briefly introduce oracle Turing machines based on [6, Section 3.4]. We say a Turing machine M^f has *oracle access to a function* $f : \{0,1\}^\star \mapsto \mathbb{N}_0$ if M^f has an additional tape and additional states f_{query}, f_0, and f_1 such that the following holds. When M^f writes $(w,(i)_2)$ on the additional tape and afterwards moves into the state f_{query}, M^f moves into the state f_1 if the ith bit of $(f(w))_2$ is 1 and otherwise into f_0, where $w \in \{0,1\}^\star$, $i \in \mathbb{N}$, and $(i)_2$ and $(f(w))_2$ denote the respective numbers base 2.

Now, we say a function $f : \{0,1\}^\star \mapsto \mathbb{N}_0$ is *#P-hard* if for all $g \in$ #P, there is a polynomial-time Turing machine M_g^f that has oracle access to f and computes g. We say f is #P-complete if $f \in$ #P and f is #P-hard. For more information on the complexity class #P, please refer to [6, Section 17.2].

10.1.3. The Complexity Classes FPT and W[1]

Usually, the running time of algorithms is analyzed as a function of the input length. In the context of *parameterized complexity*, we analyze the running time of an algorithm as a function of the input length and one or more additional parameters. The complexity class *fixed-parameter tractable problems* (FPT) captures problems for which there is an algorithm whose running time is a polynomial in the input length and a function in an additional parameter. Such problems are efficiently solvable if the additional parameter is not “too large”.

More formally, we call a relation $\mathcal{R} \subseteq \mathcal{I} \times \mathbb{N}$, where $\mathcal{I} \subseteq \{0,1\}^\star$ is the set of problem instances, a *parameterized problem*. Then, the class FPT consists of parameterized problems $\mathcal{R}$ for which there exists a computable function $f : \mathbb{N} \mapsto \mathbb{N}$, a polynomial p, and a Turing machine $M_\mathcal{R}$ such that $M_\mathcal{R}$ decides whether a pair $(I,k) \in \{0,1\}^\star \times \mathbb{N}$ is in $\mathcal{R}$ in $O\big(f(k)p(|I|)\big)$ time.

We say a parameterized problem $\mathcal{R}$ is *FPT-reducible* to a parameterized problem $\mathcal{R}'$ if there are computable functions $f, f' : \mathbb{N} \mapsto \mathbb{N}$, a polynomial p, and a function $\Phi : \{0,1\}^\star \times \mathbb{N} \mapsto \{0,1\}^\star \times \mathbb{N}$ such that

- a pair $(I,k) \in \{0,1\}^\star \times \mathbb{N}$ is in $\mathcal{R}$ if and only if $\Phi(I,k) \in \mathcal{R}'$.
- Let $(I,k) \in \{0,1\}^\star \times \mathbb{N}$ be a pair and write $\Phi(I,k) = (I',k')$. Then, we have $k' \leq f'(k)$.
- for all pairs $(I,k) \in \{0,1\}^\star \times \mathbb{N}$, $\Phi(I,k)$ can be computed in $O\big(f(k)p(|I|)\big)$ time.

There are parameterized problems that are not believed to be in FPT. One such problem is CLIQUE:

GIVEN a graph G and a parameter $k \in \mathbb{N}$,

DECIDE whether there is a clique in G of size at least k.

We define $\mathsf{W[1]}$ to be the complexity class that consists of all parameterized problems that are FPT-reducible to CLIQUE. We say a parameterized problem $\mathcal{R}$ is *$\mathsf{W[1]}$-complete* if $\mathcal{R} \in \mathsf{W[1]}$ and all parameterized problems in $\mathsf{W[1]}$ are FPT-reducible to $\mathcal{R}$. Furthermore, we say a function $g : \mathbb{N} \mapsto \mathbb{N}$ is *$\mathsf{W[1]}$-hard* if there is a $\mathsf{W[1]}$-complete parameterized problem $\mathcal{R}$ such that there exists a polynomial p, a computable function $f : \mathbb{N} \mapsto \mathbb{N}$, and an oracle Turing machine M^g with oracle access to g that decides whether a pair $(I,k) \in \{0,1\}^\star \times \mathbb{N}$ is in $\mathcal{R}$ in $O\big(f(k)p(|I|)\big)$ time.

Please refer to [25, Section 13.3] and [31, Chapter 10] for a thorough discussion of $\mathsf{W[1]}$ and the $\mathsf{W[t]}$-hierarchy, as well as for the alternative definition of $\mathsf{W[1]}$ using circuit families.

10.2. Approximation in High Dimensions

Let $P \subset \mathbb{R}^d$ be a set and $\boldsymbol{q} \in \mathbb{R}^d$ a query point. For the remainder of this section, we consider d to be constant. We distinguish to cases: if $\sigma_P(\boldsymbol{q})$ is "small", then we enumerate all $\boldsymbol{q}$-embracing $(d+1)$-subsets of P. Otherwise, if $\sigma_P(\boldsymbol{q})$ is "large", we approximate the simplicial depth by random sampling. We begin with the first case.

In the following, we assume without loss of generality that $\dim P = d$, the point $\boldsymbol{q}$ is embraced by P (since otherwise the simplicial depth is 0), and by rescaling the points in P, we can assume that $\boldsymbol{q}$ is the barycenter. Furthermore, we assume that there is no $\boldsymbol{q}$-embracing d-subset of P. Note that this assumption is NP-complete to verify [43, Theorem 5]. However, we can perturb the point $\boldsymbol{q}$ with the techniques from Chapter 3 in polynomial time such that for all $(d+1)$-subsets $S \subset P$ with $\boldsymbol{q} \in \operatorname{int}\operatorname{conv}(S)$, the perturbed point $\boldsymbol{q}^\approx$ is also embraced by S. However, if $\boldsymbol{q}$ is contained in a facet of the convex hull of some $(d+1)$-subset S', then $\boldsymbol{q}^\approx$ is not necessarily embraced by S'. Thus, the simplicial depth may change by perturbing $\boldsymbol{q}$ and it is not clear how to bound this difference.

Now, let

$$\Delta_P(\boldsymbol{q}) = \{S \subseteq P \mid |S| = d+1,\ S \text{ embraces } \boldsymbol{q}\}$$

denote the set of all $\boldsymbol{q}$-embracing $(d+1)$-subsets of P. Consider the graph $G_P(\boldsymbol{q}) = (V = \Delta_P(\boldsymbol{q}), E)$ that has one node per $\boldsymbol{q}$-embracing $(d+1)$-subset and there is an edge $\{S, S'\} \in E$ between two vertices $S, S' \in V$ if and only if S can be obtained from S' by swapping one point with a different point in P, i.e., if the symmetric difference of S and S' has size 2. Before we show that this graph is connected, we need the following simple observation.

Observation 10.3. *Let $P \subset \mathbb{R}^d$ be a point set and let $\boldsymbol{q} \in \mathbb{R}^d$ be a point such that no d-subset of P embraces $\boldsymbol{q}$. Then, for every $\boldsymbol{q}$-embracing $(d+1)$-subset $S \subseteq P$ and for every point $\boldsymbol{p} \in P$, there exists a unique point $\boldsymbol{p}' \in S$ such that the set $S' = (S \setminus \{\boldsymbol{p}'\}) \cup \{\boldsymbol{p}\}$ embraces $\boldsymbol{q}$.*

Proof. Consider the ray $\vec{r}$ that originates at $\boldsymbol{q}$ and goes in direction $\boldsymbol{q} - \boldsymbol{p}$. Since no d-subset of P embraces $\boldsymbol{q}$ and since S embraces $\boldsymbol{q}$, the ray $\vec{r}$ intersects a facet of $\operatorname{conv}(S)$ in its interior. Then, the single point $\boldsymbol{p}' \in S$ that does not define this facet is the unique point from S such that the set $(S \setminus \{\boldsymbol{p}'\}) \cup \{\boldsymbol{p}\}$ is $\boldsymbol{q}$-embracing. ∎

In the following, we say a graph $G = (V, E)$ is *k-connected* if the removal of any $k-1$ vertices does not separate G. Using Gale transform, we can show that $G_P(\boldsymbol{q})$ is isomorphic to the 1-skeleton of an $(n-d-1)$-dimensional polytope. Since Balinski's theorem [8] states that the 1-skeleton of every $(n-d-1)$-dimensional polytope is $(n-d-1)$-connected, the graph $G_P(\boldsymbol{q})$ is $(n-d-1)$-connected.

Lemma 10.4. *Let $\boldsymbol{q} \in \mathbb{R}^d$ be a point and let $P \subset \mathbb{R}^d$ be a point set of size n such that $\boldsymbol{q}$ is the barycenter of P and such that* $\dim P = d$. *Then, $G_P(\boldsymbol{q})$ is $(n-d-1)$-connected and $(n-d-1)$-regular.*

Proof. We can assume without loss of generality that $\boldsymbol{q} = \boldsymbol{0}$. Write $P = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\} \subset \mathbb{R}^d$. Since the origin is the barycenter of P, by Lemma 10.2 there exists a set $\mathrm{GT}^{-1}(P) = \{\overline{\boldsymbol{p}}_1^{-1}, \ldots, \overline{\boldsymbol{p}}_n^{-1}\} \subset \mathbb{R}^{n-d-1}$ such that $P = \mathrm{GT}(\mathrm{GT}^{-1}(P))$. Let now S, S' be two adjacent nodes in $G_P(\boldsymbol{q})$ and consider the two subsets $\overline{S} = \{\overline{\boldsymbol{p}}_i^{-1} \mid \boldsymbol{p}_i \in P \setminus S\}$ and $\overline{S}' = \{\overline{\boldsymbol{p}}_i^{-1} \mid \boldsymbol{p}_i \in P \setminus S'\}$. By Lemma 10.1, the sets $\overline{S}$ and $\overline{S}'$ define two facets f_1 and f_2 of $\operatorname{conv}(\mathrm{GT}^{-1}(P))$, respectively. Since S and S' are adjacent, we have $|S \cap S'| = d$ and hence $|\overline{S} \cap \overline{S}'| = n-d-2$. Thus, f_1 and f_2 share a ridge. Hence, $G_P(q)$ is isomorphic to the 1-skeleton of the polytope dual to $\operatorname{conv}(\mathrm{GT}^{-1}(P))$. Balinski's theorem [8] then implies that $G_P(\boldsymbol{q})$ is $(n-d-1)$-connected. Moreover, Observation 10.3 directly implies that $G_P(\boldsymbol{q})$ is $(n-d-1)$-regular. ∎

Since $G_P(\boldsymbol{q})$ is connected, the number of vertices in $G_P(\boldsymbol{q})$ can be counted by exploring the graph with breadth-first search. The following result is now immediate.

Lemma 10.5. *Let $P \subset \mathbb{R}^d$ be a point set and let $\boldsymbol{q} \in \mathbb{R}^d$ be a point. Then, $\sigma_P(\boldsymbol{q})$ can be computed in $O(n\sigma_P(\boldsymbol{q}))$ time.*

If the simplicial depth is large, the enumeration approach becomes infeasible. In this case, we apply a simple random sampling algorithm.

Lemma 10.6. *Let $P \subset \mathbb{R}^d$ be a set and let $\boldsymbol{q} \in \mathbb{R}^d$ be a point. Furthermore, let $\varepsilon, \delta > 0$ be constants and let $m \in \mathbb{N}$ be a parameter. If $\sigma_P(\boldsymbol{q}) \geq m$, then $\sigma_P(\boldsymbol{q})$ can be $(1+\varepsilon)$-approximated in $\widetilde{O}\left(n^{d+1}/m\right)$ time with error probability $O\left(n^{-\delta}\right)$.*

Proof. Let $S_1, \ldots, S_k$ be k random $(d+1)$-subsets of P for

$$k = \left\lceil \binom{n}{d+1} \frac{4\delta \log n}{\varepsilon^2 m} \right\rceil.$$

For each random subset S_i, let X_i be 1 if S_i embraces $\boldsymbol{q}$ and 0 otherwise. Then, we have

$$\mu = \mathbb{E}\left[\sum_{i=1}^{k} X_i\right] = k \frac{\sigma_P(q)}{\binom{n}{d+1}} \geq \frac{4\delta\sigma_P(q)}{\varepsilon^2 m} \log n \geq \frac{4\delta}{\varepsilon^2} \log n.$$

Applying the Chernoff bound [6, Corollorary A.15], we get

$$\Pr\left[\left|\sum_{i=1}^{k} X_i - \mu\right| \geq \varepsilon\mu\right] \leq 2\exp\left(-\frac{\varepsilon^2}{4}\mu\right) \leq 2n^{-\delta}.$$

Thus, $\binom{n}{d+1} k^{-1} \sum_{i=1}^{k} X_i$ is a $(1+\varepsilon)$-approximation of $\sigma_P(q)$ with error probability $O\left(n^{-\delta}\right)$.

For $d = O(1)$, we can test in $O(1)$ whether a given $(d+1)$-subset of P embraces a point. Hence, the running time is dominated by the number of samples. ■

Now, combining both algorithms from Lemma 10.5 and from Lemma 10.6 enables us to approximate the simplicial depth in $\widetilde{O}\left(n^{d/2+1}\right)$ time.

Theorem 10.7. *Let $P \subset \mathbb{R}^d$ and let $\boldsymbol{q} \in \mathbb{R}^d$ be a point. Furthermore, let $\varepsilon > 0$ and $\delta > 0$ be constants. Then, $\sigma_P(\boldsymbol{q})$ can be $(1+\varepsilon)$-approximated in $\widetilde{O}\left(n^{d/2+1}\right)$ time with error probability $O\left(n^{-\delta}\right)$.*

Proof. We apply the algorithm from Lemma 10.5 and stop it once $m = \lceil n^{d/2} \rceil$ nodes of $G_P(\boldsymbol{q})$ are explored. This requires $O\left(n^{d/2+1}\right)$ time. If the graph is fully explored, we know the value of $\sigma_P(\boldsymbol{q})$ and return it. Otherwise, the simplicial depth $\sigma_P(q)$ is at least m. We can now apply the algorithm from Lemma 10.6 and compute a $(1+\varepsilon)$-approximation in $\widetilde{O}\left(n^{d/2+1}\right)$ time with error probability $O\left(n^{-\delta}\right)$. ■

10.3. Improving the Brute-Force Approach

We now consider the problem of computing the simplicial depth exactly. Again, we assume that the dimension is constant. The algorithm improves on the naive $\Omega\left(n^{d+1}\right)$ algorithm that tests every $(d+1)$-subset of the input point set by using the following simplex range searching data structure from Chazelle et al. [18].

Theorem 10.8 ([18, Theorem 2.1]). *Let $P \subset \mathbb{R}^d$ be a set of n points and let $\varepsilon > 0$ be an arbitrary constant. Then there exists a data structure that, given a query simplex $\sigma \subseteq \mathbb{R}^d$, supports two type of queries: (i) it counts the number of points s in $P \cap \sigma$ in $O\left(\log n\right)$ time, and (ii) it reports* *the points in $P \cap \sigma$ in $O(\log n + s)$ time. Furthermore, this data structure can be constructed with $O(n^{d+\varepsilon})$ preprocessing time and space.* ■

Let now $P \subset \mathbb{R}^d$ be a point set and $\boldsymbol{q} \in \mathbb{R}^d$ for which we want to compute the simplicial depth. Let h denote the hyperplane $\{\boldsymbol{x} \in \mathbb{R}^d \mid (\boldsymbol{x})_d = 1\}$ and let h' denote the parallel hyperplane $\{\boldsymbol{x} \in \mathbb{R}^d \mid (\boldsymbol{x})_d = -1\}$. Without loss of generality, we may assume that $\boldsymbol{q} = \boldsymbol{0}$ and that each ray $\operatorname{pos}(\boldsymbol{p})$, $\boldsymbol{p} \in P$, intersects one of the two hyperplanes. Now, we enumerate all possible d-subsets $P' = \{\boldsymbol{p}'_1, \ldots, \boldsymbol{p}'_d\}$ of P and for each such set, we count the number of points in P that extend P' to a $\boldsymbol{0}$-embracing $(d+1)$-subset of P as follows. First, we observe that all such points $\boldsymbol{p}' \in P$ must be contained in the cone $\operatorname{pos}\{-\boldsymbol{p}'_1, \ldots, -\boldsymbol{p}'_d\}$ and in particular, we must have either

$$(h \cap \operatorname{pos}(\boldsymbol{p}')) \subset \sigma_h = h \cap \operatorname{pos}\{-\boldsymbol{p}'_1, \ldots, -\boldsymbol{p}'_d\},$$

or

$$(h' \cap \operatorname{pos}(\boldsymbol{p}')) \subset \sigma_{h'} = h' \cap \operatorname{pos}\{-\boldsymbol{p}'_1, \ldots, -\boldsymbol{p}'_d\}.$$

Let $P_h = \{h \cap \operatorname{pos}(\boldsymbol{p}) \mid \boldsymbol{p} \in P\}$ denote the intersections of the rays through the points in P with h and similarly, let $P_{h'} = \{h' \cap \operatorname{pos}(\boldsymbol{p}) \mid \boldsymbol{p} \in P\}$ denote the intersections with h'. We construct one data structure D_h from Theorem 10.8 for P_h and one data structure $D_{h'}$ for $P_{h'}$. Now, we can count the number of points in $\sigma_h \cap P_h$ and the number of points in $\sigma_{h'} \cap P_{h'}$ with the data structures D_h and $D_{h'}$, respectively. Please refer to Figure 10.2 (a) and (b) for an example. Using this approach, we count each $\boldsymbol{0}$-embracing $(d+1)$-subset $d+1$ times, one time for each d-subset. Hence, the total sum divided by $d+1$ is the simplicial depth of the origin with respect to P. See Algorithm 10.1 for details.

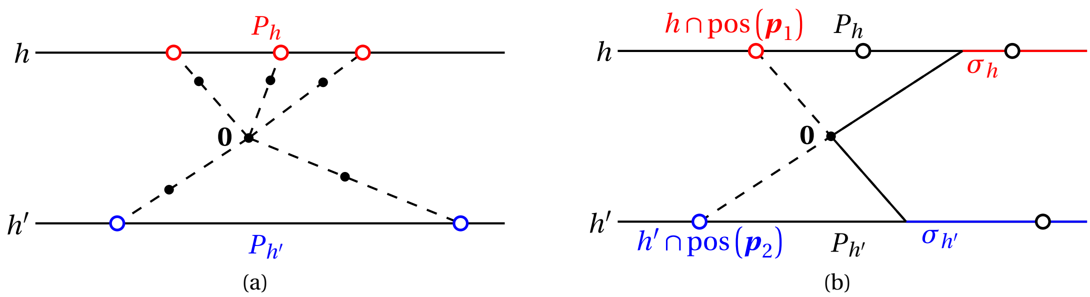

Figure 10.2.: (a) The projection of the points from P onto the two hyperplanes h and h'. (b) To count the number of points that extend $\{\boldsymbol{p}_1, \boldsymbol{p}_2\}$ to a $\boldsymbol{0}$-embracing 3-set, the two data structures are queried with the simplices σ_h and $\sigma_{h'}$.

It remains to analyze the running time of Algorithm 10.1. Since we construct the two data structures for a $(d-1)$-dimensional subspace, this step requires $O(n^{d-1+\varepsilon})$ time, where $\varepsilon > 0$ is an arbitrarily small constant. We query the two data structures for each of the $\Theta(n^d)$ choices for d-subsets of P and each query needs $O(\log n)$ time. Hence, the total running time is $O(n^{d-1+\varepsilon} + n^d \log n) = O(n^d \log n)$ time for $\varepsilon \leq 1$. The following theorem is now immediate.

Theorem 10.9. *Let $P \subset \mathbb{R}^d$ be a set of size n and let $\boldsymbol{q} \in \mathbb{R}^d$ be a point. Then, Algorithm 10.1 computes the simplicial depth $\sigma_P(\boldsymbol{q})$ of $\boldsymbol{q}$ with respect to P in $O(n^d \log n)$ time.*

Algorithm 10.1: Improved brute-force algorithm to compute the simplicial depth

Input: point set $P \subset \mathbb{R}^d$

$h \leftarrow \{\boldsymbol{x} \in \mathbb{R}^d \mid (\boldsymbol{x})_d = 1\}$; $h' \leftarrow \{\boldsymbol{x} \in \mathbb{R}^d \mid (\boldsymbol{x})_d = -1\}$;

$D_h \leftarrow$ data structure from Theorem 10.8 for the set $P_h = \{h \cap \operatorname{pos}(\boldsymbol{p}) \mid \boldsymbol{p} \in P\}$;

$D_{h'} \leftarrow$ data structure from Theorem 10.8 for the set $P_{h'} = \{h' \cap \operatorname{pos}(\boldsymbol{p}) \mid \boldsymbol{p} \in P\}$;

$s \leftarrow 0$;

foreach *d-subset $P' = \{\boldsymbol{p}'_1, \ldots, \boldsymbol{p}'_d\}$ of P* **do**

  $\sigma_h = h \cap \operatorname{pos}\{-\boldsymbol{p}'_1, \ldots, -\boldsymbol{p}'_d\}$; $\sigma_{h'} = h' \cap \operatorname{pos}\{-\boldsymbol{p}'_1, \ldots, -\boldsymbol{p}'_d\}$;

  Count the number of points $n_h = |P_h \cap \sigma_h|$ with D_h;

  Count the number of points $n_{h'} = |P_{h'} \cap \sigma_{h'}|$ with $D_{h'}$;

  $s \leftarrow s + n_h + n_{h'}$;

return $s/(d+1)$;

10.4. Complexity

If the dimension is constant, even the naive algorithm that checks all possible simplices needs only polynomial time. We now consider the case that d is part of the input and we show that then computing the simplicial depth is $\#\mathsf{P}$-complete by a reduction from counting the number of perfect matchings in bipartite graphs. Furthermore, by a different reduction, we show that computing the simplicial depth is $\mathsf{W}[1]$-hard with respect to the dimension as parameter.

Theorem 10.10. *Let $P \subset \mathbb{Q}^d$ be a set and $\boldsymbol{q} \in \mathbb{Q}^d$ a point. Then, computing $\sigma_P(\boldsymbol{q})$ is $\#\mathsf{P}$-complete if the dimension is part of the input.*

Proof. For a set of points $P' \subset \mathbb{Q}^d$, let $\widetilde{P}' \in \{0,1\}^\star$ denote some fixed encoding as a binary string. Similarly, we denote for a point $\boldsymbol{p} \in \mathbb{Q}^d$ with $\tilde{\boldsymbol{p}} \in \{0,1\}^\star$ an encoding of $\boldsymbol{p}$ as a binary string. Now, let $P \subset \mathbb{Q}^d$ be a set and $\boldsymbol{q} \in \mathbb{Q}^d$ be a point, and let $f_\sigma : \{0,1\}^\star \mapsto \mathbb{N}$ be the function that maps the pair $(\widetilde{P}, \tilde{\boldsymbol{q}})$ to $\sigma_P(\boldsymbol{q})$. If the argument of f_σ does not encode a pair $(P, \boldsymbol{q})$, the value of f_σ is defined as 0. It is easy to see that $f_\sigma \in \#\mathsf{P}$. Let M be the Turing machine that on input $(\widetilde{P}, \tilde{\boldsymbol{q}}, \widetilde{P}')$ checks whether $\tilde{\boldsymbol{q}}$ encodes a point $\boldsymbol{q} \in \mathbb{Q}^d$ and whether $\widetilde{P}$ and $\widetilde{P}'$ encode two sets P and $P' \subset \mathbb{Q}^d$, respectively. Furthermore, M checks whether P' is a $(d+1)$-subset of P. If one of the checks fails, M returns NO. Otherwise, M solves a linear program to decide whether $\boldsymbol{q} \in \operatorname{conv}(P')$ and it returns YES if and only if $\boldsymbol{q} \in \operatorname{conv}(P')$ and otherwise NO. Then, $f_\sigma(\widetilde{P}, \tilde{\boldsymbol{q}})$ is the number of binary strings $\widetilde{P}' \in \{0,1\}^\star$ such that M returns on input $(\widetilde{P}, \tilde{\boldsymbol{q}}, \widetilde{P}')$ YES, and M can be defined such that it returns YES or NO in polynomial time.

Let $G = (V, E)$ be a bipartite graph with $|V| = n$ and $|E| = m$. It is well known that computing the number of perfect matchings in G is $\#\mathsf{P}$-complete [79, Theorem 1]. Let $\mathcal{P}_H \subset \mathbb{R}^m$ be the perfect matching polytope for G [34, Chapter 30]. The polytope $\mathcal{P}_H$ is an H-polytope that is defined by $m + 2n$ half-spaces and its number of vertices of $\mathcal{P}_H$ equals the number k of perfect matchings in G. Consider the dual V-polytope $\mathcal{P}_V \subset \mathbb{R}^m$. Now, $\mathcal{P}_V$ is the convex hull of $m + 2n$ points $P \subset \mathbb{Q}^m$ and the number of facets of $\mathcal{P}_V$ is k. Let $\operatorname{GT}(P) \subset \mathbb{R}^{2n-1}$ be the Gale transform of P. By Lemma 10.1, there is a bijection between the facets of $\mathcal{P}_V$ and the $(2n-1)$-simplices with vertices in $\mathrm{GT}(P)$ that embrace the origin. Hence, the simplicial depth $\sigma_{\mathrm{GT}(P)}(\mathbf{0})$ of the origin with respect to $\mathrm{GT}(P)$ equals the number k of perfect matchings in G. Since all steps can be carried out in polynomial time and since the number of perfect matchings equals $\sigma_{\mathrm{GT}(P)}(\mathbf{0})$, i.e. the reduction is *parsimonious*, the function f_σ is #P-complete. ∎

By a reduction to *d-Carathéodory*, we further show that computing the simplicial depth is W[1]-hard with respect to the parameter d. In d-Carathéodory, we are given a set $P \subset \mathbb{Q}^d$ and we want to decide whether there is a $(d-1)$-simplex with vertices in P that embraces the origin. Knauer et al. [43, Theorem 5] proved that this problem is W[1]-hard with respect to the parameter d.

Theorem 10.11. *Let $P \subset \mathbb{Q}^d$ be a set and $\boldsymbol{q} \in \mathbb{Q}^d$ be a point. Then, computing $\sigma_P(\boldsymbol{q})$ is W[1]-hard with respect to the parameter d.*

Proof. Without loss of generality let $\boldsymbol{q}$ be the origin. Assume we have access to an oracle that, given a query point $\boldsymbol{q}' \in \mathbb{Q}^d$ and a set $Q \subset \mathbb{Q}^d$, returns the simplicial depth $\sigma_Q(\boldsymbol{q}')$ of $\boldsymbol{q}'$ with respect to Q. We show that then d-Carathéodory can be decided by only two oracle queries.

For $k \in \mathbb{N}$, let Δ_k denote the family of sets

$$\Delta_k = \{S \subseteq P \mid S \text{ embraces } \mathbf{0}, |S| = k-1\}$$

and let δ_k denote its size. Then, $\sigma_P(\mathbf{0}) = \delta_d$ and we want to decide whether $\delta_{d-1} > 0$. Now, for each point $\boldsymbol{p} \in P$, we denote with

$$\hat{\boldsymbol{p}} = \begin{pmatrix} \boldsymbol{p} \\ 1 \end{pmatrix} \in \mathbb{Q}^{d+1}$$

the $(d+1)$-dimensional point that we obtain by appending a 1-coordinate and we set $\widehat{P} = \{\hat{\boldsymbol{p}} \mid \boldsymbol{p} \in P\}$. Set $Q' = \{(0,\dots,0,-1)^T, (0,\dots,0,-2)^T\} \subset \mathbb{Q}^{d+1}$ and set $Q = \widehat{P} \cup Q'$. The configuration is depicted in Figure 10.3 (a). We can relate $\sigma_Q(\mathbf{0})$ with $\sigma_P(\mathbf{0})$ as follows. First, we observe that each $\mathbf{0}$-embracing $(d+2)$-subset of Q corresponds to a $\mathbf{0}$-embracing subset of P: let $\widehat{S} = \widehat{S}_P \cup \widehat{S}_{Q'} \subset Q$ be a $\mathbf{0}$-embracing set of size $|\widehat{S}| = d+2$, where $\widehat{S}_P \subseteq \widehat{P}$ and $\widehat{S}_{Q'} \subseteq Q'$. By construction of Q, $\widehat{S}_P$ must embrace $\boldsymbol{p_0} = (0,\dots,0,1)^T$ and hence the set $\{\boldsymbol{p} \mid \hat{\boldsymbol{p}} \in \widehat{S}_P\} \subseteq P$ is $\mathbf{0}$-embracing. Now, each $\boldsymbol{p_0}$-embracing set $\widehat{S}_P \subset \widehat{P}$ of size $d+1$ can be extended in two ways to a $\mathbf{0}$-embracing $(d+2)$-set of Q by taking either point from Q'. Furthermore, each $\boldsymbol{p_0}$-embracing set $\widehat{S}_P \subset \widehat{P}$ of size d can be extended in at least one way to a $\mathbf{0}$-embracing $(d+2)$-set of Q by taking both points from Q'. Hence, we obtain

$$\sigma_Q(\mathbf{0}) \begin{cases} = 2\delta_d \text{ if } \delta_{d-1} = 0, \text{ and} \\ > 2\delta_d \text{ otherwise.} \end{cases}$$

Thus, $\delta_{d-1} > 0$ if and only if $\sigma_Q(\mathbf{0}) - 2\sigma_P(\mathbf{0}) > 0$. Since the construction of Q takes only polynomial time, the problem of computing the simplicial depth is W[1]-hard with respect to the parameter d. ∎

As discussed at the beginning of this chapter, Theorem 10.11 implies the following corollary.

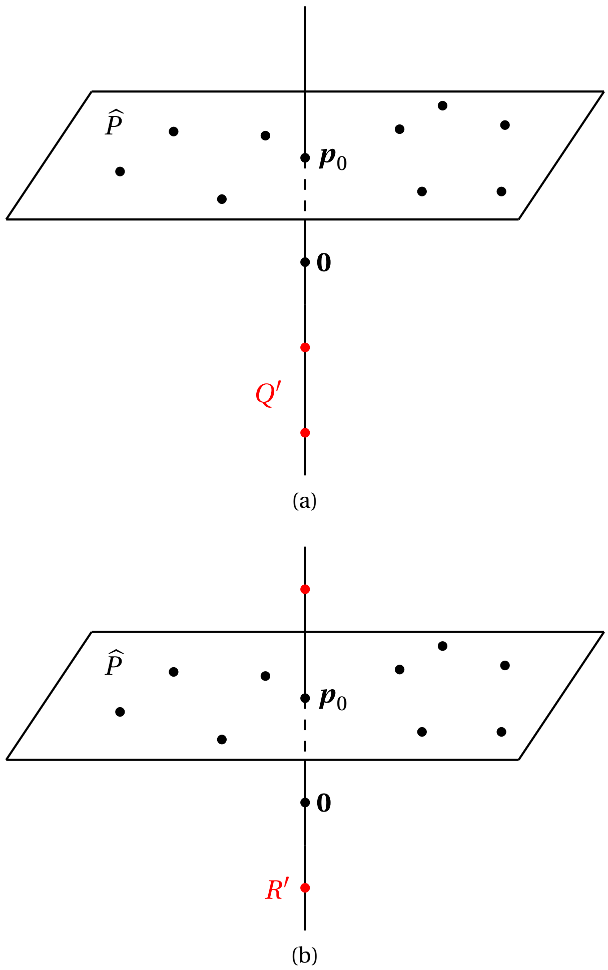

Figure 10.3.: (a) The point set Q consists of the lifted points $\widehat{P}$ and the two additional points from Q' that each can extend a $\boldsymbol{p_0}$-embracing $(d+1)$-subset of $\widehat{P}$ to a $\mathbf{0}$-embracing $(d+2)$-subset of Q. (b) The point set R consists of the lifted points $\widehat{P}$ and the two additional points from R'. Only one point from R' can extend a $\boldsymbol{p_0}$-embracing $(d+1)$-subset of $\widehat{P}$ to a $\mathbf{0}$-embracing $(d+2)$-subset of R.

Corollary 10.12. *Let $C_1, \dots, C_m \subset \mathbb{Q}^d$ be m color classes and let $\boldsymbol{q} \in \mathbb{Q}^d$ be a point. Then, computing the colorful simplicial depth $\sigma^{col}_{C_1,\dots,C_m}(\boldsymbol{q})$ of $\boldsymbol{q}$ with respect to $C_1, \dots, C_m$ is W[1]-hard with respect to the parameter d.*

We conclude this chapter by showing that although computing the simplicial depth is #P-complete, it is possible to determine the parity in polynomial time.

Theorem 10.13. *Let $P \subset \mathbb{R}^d$ be a set of n points and let $\boldsymbol{q} \in \mathbb{R}^d$ be a point such that there is no $\boldsymbol{q}$-embracing d-subset of P. Then $\sigma_P(\boldsymbol{q})$ is even if and only if $n-d-1$ is odd or $\binom{n}{d}$ is even. Equivalently, $\sigma_P(\boldsymbol{q})$ is even if and only if $n-d-1$ is odd or there exists an $i \in \mathbb{N}$ such that the ith bit of $(n-d)_2$ and $(d)_2$ is 1, where $(n-d)_2$ and $(d)_2$ denote the numbers n and d written in binary, respectively.*

Proof. We assume without loss of generality that $\boldsymbol{q}$ is the origin. Since there is no d-set $P' \subseteq P$ that embraces $\boldsymbol{q}$, Lemma 10.4 states that the graph $G_P(\mathbf{0})$ is $(n-d-1)$-regular. Then, the product $(n-d-1)|V| = (n-d-1)\sigma_P(\mathbf{0})$ is even. If $(n-d-1)$ is odd, $\sigma_P(\boldsymbol{q})$ must be even. Assume now $(n-d-1)$ is even. We construct a new point set R in $\mathbb{R}^{d+1}$ similar as in the proof of Theorem 10.11: let R' denote the set $\{(0,\dots,0,-1)^T, (0,\dots,0,2)^T\} \subset \mathbb{R}^{d+1}$ and set $R = \widehat{P} \cup R' \subset \mathbb{R}^{d+1}$, where $\widehat{P}$ is defined as in the proof of Theorem 10.11. See Figure 10.3 (b). Let us now consider the graph $G_R(\mathbf{0})$. Since $n-d-1$ is even, $(|R|-(d+1)-1) = n-d$ is odd. Because there is no d-subset of P that embraces the origin, there is no $(d+1)$-subset of R that embraces the origin. Then, Lemma 10.4 implies that $G_R(\mathbf{0})$ is $(n-d)$-regular and thus $\sigma_R(\mathbf{0})$ is even. Let $\widehat{S} \subset R$, $|R| = d+2$, be a subset that contains the origin in its convex hull. Then either (i) $R' \subset \widehat{S}$ or (ii) $\widehat{S}$ contains the point $\boldsymbol{r} = (0,\dots,0,-1)^T \in R'$ and $d+1$ points $\widehat{S}_P \subseteq \widehat{P}$ with $(0,\dots,0,1)^T \in \operatorname{conv}(\widehat{S}_P)$. There are $\binom{n}{d}$ sets $\widehat{S}$ with Property (i) and $\sigma_P(\mathbf{0})$ sets $\widehat{S}$ with Property (ii). Hence, we have $\sigma_R(\mathbf{0}) = \sigma_P(\mathbf{0}) + \binom{n}{d}$ is even and thus $\sigma_P(\mathbf{0})$ is odd if and only if $\binom{n}{d}$ is odd.

Finally, we can use Kummer's theorem [35,45] to obtain a precise criterion when $\binom{n}{d}$ is even. Let $p \in \mathbb{N}$ be a prime. Then, Kummer's theorem states that the maximum power $k \in \mathbb{N}$ such that p^k divides $\binom{n}{d}$ is the number of carries when we add $(n-d)_p$ and $(d)_p$, where $(n-d)_p$ and $(d)_p$ denote the numbers $n-d$ and d written in base p, respectively. In particular, p divides $\binom{n}{d}$ if there is at least one carry. Thus, for $p=2$, the binomial coefficient is divisible by 2 and hence even if and only if there is an index $i \in \mathbb{N}$ such that the ith bit of the binary numbers $(n-d)_2$ and $(d)_2$ is 1. ∎

11 Conclusions

We have seen new upper bounds on the complexity of ColorfulCarathéodory and its descendants, Centerpoint, Tverberg, SimplicialCenter, and ColorfulKirchberger that put all five problems on the edge of the three complexity classes FP, PPAD, and PLS. Now, showing that one of the above problems is PPAD- or PLS-complete would imply PPAD $\subseteq$ PLS or PLS $\subseteq$ PPAD, and a polynomial-time algorithm would be interesting for its own merit. We believe that these results are only a first step towards settling the complexity of ColorfulCarathéodory and its descendants and that pursuing this question will help us to improve our understanding of the relationship between FP, PPAD, and PLS. We conclude with open problems.

The complexity of ColorfulCarathéodory. The intersection of PPAD and PLS contains the complexity class *continuous local search* (CLS), which in turn then contains FP. Clearly, the most pressing open question is whether ColorfulCarathéodory is contained in CLS. Neither the reduction of ColorfulCarathéodory to Sperner nor the formulation of ColorfulCarathéodory as a PLS-problem can be extended directly to a formulation of ColorfulCarathéodory as a CLS-problem and it seems that a fundamentally different approach would be necessary to achieve this.

Approximating ColorfulCarathéodory. The dimension-reduction argument that has been presented in Chapter 7 leads to an algorithm that computes $\lceil \varepsilon d \rceil$-colorful choices in polynomial time for any fixed $\varepsilon > 0$. Is it possible to compute a $\sqrt{d}$-colorful choice in polynomial time? On the other hand, is it possible to show that computing a 2-colorful choice is as hard as solving ColorfulCarathéodory exactly? Furthermore, we have seen in Chapter 8 that given two $\mathbf{0}$-embracing color classes $C_1, C_2 \subset \mathbb{Q}^d$ that ray-embrace a point $\boldsymbol{b} \in \mathbb{Q}^d$, $\boldsymbol{b} \neq \mathbf{0}$, and a number $k \in [d-1]$, it is possible to compute a $\mathbf{0}$-embracing $(k, d-k)$-colorful choice in weakly polynomial time. In the algorithm, we used a binary search approach by reducing ColorfulCarathéodory with two colors to Sperner's lemma in one dimension. In general, we can reduce ColorfulCarathéodory with m colors to Sperner's lemma in $m-1$ dimensions. However, for $m > 2$ it is no longer clear how to determine a part of the simplex that contains a solution. Can the parameter space and or the cost function in the reduction from ColorfulCarathéodory to Sperner's lemma be modified to support such a recursive approach?

Exact solutions for ColorfulCarathéodory. In Chapter 9, we discussed the problem of solving ColorfulCarathéodory exactly. Can the recent $O(n^{d-1})$ algorithm for computing the simplicial depth by Pilz et al. [62] be adapted to improve the naive algorithm for Col-

ORFULCARATHÉODORY? Furthermore, we have seen that with $\Theta\left(d^2 \log d\right)$ color classes, the problem can be solved in quasi-polynomial time. Is it possible to solve the exact problem in polynomial-time when given poly(d) many color classes?

Approximating the Simplicial Depth. The approximation algorithm in Section 10.2 combines two different strategies: first, we traverse a graph with one node per simplex that contains the query point until up to $\lceil n^{d/2} \rceil$ distinct nodes have been discovered. If the graph is fully explored, we know the exact simplicial depth and otherwise we approximate the simplicial depth with random sampling. Exploring the graph with BFS requires $\Omega\left(n^{d/2}\right)$ space when implemented naively. Can the need for space be improved by a careful enumeration of the nodes in the graph? Furthermore, although we know by Lemma 10.4 that the constructed graph is highly structured, we are only using the fact that it is connected. It seems plausible that the high connectedness and the regularity can be further algorithmically exploited.

Bibliography

[1] E. Aarts and J. K. Lenstra. *Local search in combinatorial optimization.* Princeton University Press, 2003. (refs: pp. 23, 24)

[2] P. Afshani, D. R. Sheehy, and Y. Stein. Approximating the simplicial depth. arXiv:1512.04856, 2015. (ref: p. 15)

[3] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The design and analysis of computer algorithms.* Addison-Wesley, 1974. (ref: p. 2)

[4] K. M. Anstreicher and R. A. Bosch. Long steps in an $O\left(n^3 L\right)$ algorithm for linear programming. *Math. Program.*, 54(1-3):251–265, 1992. (refs: pp. 3, 15)

[5] J. L. Arocha, I. Bárány, J. Bracho, R. Fabila, and L. Montejano. Very colorful theorems. *Discrete Comput. Geom.*, 42(2):142–154, 2009. (refs: pp. 6, 10, 12)

[6] S. Arora and B. Barak. *Computational complexity: a modern approach.* Cambridge University Press, 2009. (refs: pp. 71, 98, 101)

[7] F. Aurenhammer. Using Gale transforms in computational geometry. *Math. Program.*, 52(1):179–190, 1991. (ref: p. 97)

[8] M. L. Balinski. On the graph structure of convex polyhedra in n-space. *Pacific J. Math.*, 11:431–434, 1961. (ref: p. 100)

[9] I. Bárány. A generalization of Carathéodory's theorem. *Discrete Math.*, 40(2-3):141–152, 1982. (refs: pp. 4, 6, 13, 14, 60, 91)

[10] I. Bárány and J. Matoušek. Quadratically many colorful simplices. *SIAM J. Discrete Math.*, 21(1):191–198, 2007. (ref: p. 7)

[11] I. Bárány and S. Onn. Colourful linear programming and its relatives. *Math. Oper. Res.*, 22(3):550–567, 1997. (refs: pp. 6, 10, 21, 60, 65, 69, 95)

[12] 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. (ref: p. 69)

[13] A. Blum, S. Har-Peled, and B. Raichel. Sparse approximation via generating point sets. In *Proc. 27th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 548–557, 2016. (ref: p. 4)

[14] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan. Time bounds for selection. *J. Comput. System Sci.*, 7(4):448–461, 1973. (ref: p. 10)

[15] B. Bukh, J. Matoušek, and G. Nivasch. Stabbing simplices by points and flats. *Discrete Comput. Geom.*, 43(2):321–338, 2008. (ref: p. 14)

[16] C. Carathéodory. Über den Variabilitätsbereich der Fourier'schen Konstanten von positiven harmonischen Funktionen. *Rendiconti del Circolo Matematico di Palermo*, 32(1):193–217, 1911. (ref: p. 2)

[17] T. M. Chan. An optimal randomized algorithm for maximum Tukey depth. In *Proc. 15th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 430–436, 2004. (refs: pp. 1, 10)

[18] B. Chazelle, M. Sharir, and E. Welzl. Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Algorithmica*, 8(5-6):407–429, 1992. (ref: p. 101)

[19] X. Chen and X. Deng. 3-Nash is PPAD-complete. In *Electron. Colloq. Comput. Complex. (ECCC)*, volume 134, 2005. (ref: p. 23)

[20] X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *Proc. 47th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 261–272, 2006. (ref: p. 23)

[21] X. Chen and X. Deng. On the complexity of 2D discrete fixed point problem. *Theoret. Comput. Sci.*, 410(44):4448–4456, 2009. (ref: p. 23)

[22] A. Y. Cheng and M. Ouyang. On algorithms for simplicial depth. In *Proc. 13th Canad. Conf. Comput. Geom. (CCCG)*, pages 53–56, 2001. (ref: p. 95)

[23] V. Chvátal. *Linear programming.* W. H. Freeman and Company, 1983. (ref: p. 25)

[24] D. I. A. Cohen. On the Sperner lemma. *J. Combinatorial Theory*, 2(4):585–587, 1967. (refs: pp. 21, 35)

[25] M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized algorithms.* Springer, 2015. (refs: pp. 71, 99)

[26] G. B. Dantzig. *Linear Programming and Extensions.* Princeton University Press, 1963. (ref: p. 29)

[27] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009. (ref: p. 23)

[28] C. Daskalakis and C. Papadimitriou. Continuous local search. In *Proc. 22nd Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 790–804, 2011. (ref: p. 21)

[29] A. Deza, S. Huang, T. Stephen, and T. Terlaky. Colourful simplicial depth. *Discrete Comput. Geom.*, 35(4):597–615, 2006. (ref: p. 6)

[30] A. Deza, F. Meunier, and P. Sarrabezolles. A combinatorial approach to colourful simplicial depth. *SIAM J. Discrete Math.*, 28(1):306–322, 2014. (ref: p. 7)

[31] R. G. Downey and M. R. Fellows. *Parameterized complexity*. Springer, 1999. (refs: pp. 71, 99)

[32] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. 36th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 604–612, 2004. (ref: p. 25)

[33] J. Gil, W. Steiger, and A. Wigderson. Geometric medians. *Discrete Math.*, 108(1):37–51, 1992. (ref: p. 95)

[34] R. L. Graham, M. Grötschel, and L. Lovász. *Handbook of combinatorics*, volume 2. Elsevier, 1995. (ref: p. 103)

[35] A. Granville. Zaphod Beeblebrox's brain and the fifty-ninth row of Pascal's triangle. *Amer. Math. Monthly*, 99(4):318–331, 1992. (ref: p. 106)

[36] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 32:175–176, 1923. (ref: p. 7)

[37] A. F. Holmsen, J. Pach, and H. Tverberg. Points surrounding the origin. *Combinatorica*, 28(6):633–644, 2008. (ref: p. 6)

[38] S. Jadhav and A. Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. *Discrete Comput. Geom.*, 12(3):291–312, 1994. (ref: p. 10)

[39] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Comput. System Sci.*, 37(1):79–100, 1988. (refs: pp. 20, 23, 24, 57)

[40] 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. (refs: pp. 57, 59, 69)

[41] S. Khuller and J. S. B. Mitchell. On a triangle counting problem. *Inf. Process. Lett.*, 33(6):319–321, 1990. (ref: p. 95)

[42] P. Kirchberger. Über Tchebychefsche Annäherungsmethoden. *Math. Ann.*, 57(4):509–540, 1903. (ref: p. 12)

[43] C. Knauer, H. R. Tiwary, and D. Werner. On the computational complexity of ham-sandwich cuts, Helly sets, and related problems. In *Proc. 28th Int. Sympos. Theoret. Aspects of Comput. Sci. (STACS)*, volume 9, pages 649–660, 2011. (refs: pp. 4, 99, 104)

[44] 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. (refs: pp. 57, 59, 69)

[45] E. E. Kummer. Über die Ergänzungssätze zu den allgemeinen Reciprocitätsgesetzen. *J. Reine Angew. Math.*, 44:93–146, 1852. (ref: p. 106)

[46] R. Y. Liu. On a notion of data depth based on random simplices. *Ann. Statist.*, 18(1):405–414, 1990. (ref: p. 13)

[47] L. Lovász. Problem 206. *Matematikai Lapok*, 25:181, 1974. (ref: p. 7)

[48] J. Matoušek. *Lectures on discrete geometry*. Springer, 2002. (refs: pp. 7, 9, 13, 95)

[49] J. Matoušek. *Using the Borsuk-Ulam theorem: lectures on topological methods in combinatorics and geometry*. Springer, 2008. (ref: p. 45)

[50] J. Matoušek and B. Gärtner. *Understanding and using linear programming*. Springer, 2007. (ref: p. 27)

[51] N. Megiddo and R. Chandrasekaran. On the ϵ-perturbation method for avoiding degeneracy. *Oper. Res. Lett.*, 8(6):305–308, 1989. (ref: p. 29)

[52] N. Megiddo and C. H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theoret. Comput. Sci.*, 81(2):317–324, 1991. (ref: p. 20)

[53] F. Meunier and A. Deza. A further generalization of the colourful Carathéodory theorem. In *Discrete geometry and optimization*, volume 69, pages 179–190. 2013. (refs: pp. 6, 21)

[54] F. Meunier, W. Mulzer, P. Sarrabezolles, and Y. Stein. The rainbow at the end of the line — a PPAD formulation of the colorful Carathéodory theorem with applications. In *Proc. 28th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, 2017. To appear. (ref: p. 15)

[55] F. Meunier and P. Sarrabezolles. Colorful linear programming, Nash equilibrium, and pivots. arXiv:1409.3436, 2014. (refs: pp. 19, 21, 60, 65)

[56] W. Michiels, E. Aarts, and J. Korst. *Theoretical aspects of local search*. Springer, 2007. (ref: p. 23)

[57] G. L. Miller and D. R. Sheehy. Approximate centerpoints with proofs. *Comput. Geom.*, 43(8):647–654, 2010. (refs: pp. 10, 91, 92)

[58] 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. (ref: p. 15)

[59] W. Mulzer and D. Werner. Approximating Tverberg points in linear time for any fixed dimension. *Discrete Comput. Geom.*, 50(2):520–535, 2013. (ref: p. 10)

[60] C. H. Papadimitriou. The complexity of the Lin-Kernighan heuristic for the traveling salesman problem. *SIAM J. Comput.*, 21(3):450–465, 1992. (ref: p. 25)

[61] 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. (refs: pp. 20, 22, 23, 38, 47, 51, 52)

[62] A. Pilz, E. Welzl, and M. Wettstein. From crossing-free graphs on wheel sets to embracing simplices and polytopes with few vertices. Manuscript, 2016. (refs: pp. 95, 107)

[63] F. P. Preparata and M. I. Shamos. *Computational geometry.* Springer, 1985. (ref: p. 2)

[64] H. Rademacher and I. J. Schoenberg. Helly's theorem on convex domains and Tchebycheff's approximation problem. *Canad. J. Math.*, 2:245–256, 1950. (ref: p. 12)

[65] R. Rado. A theorem on general measure. *J. London Math. Soc.*, 21:291–300 (1947), 1946. (refs: pp. 1, 9)

[66] D. Rolnick and P. Soberón. Algorithmic aspects of Tverberg's theorem. arXiv:1601.03083, 2016. (ref: p. 14)

[67] J.-P. Roudneff. Partitions of points into simplices with k-dimensional intersection. I. The conic Tverberg's theorem. *European J. Combin.*, 22(5):733–743, 2001. (ref: p. 10)

[68] P. J. Rousseeuw and I. Ruts. Bivariate location depth. *J. R. Stat. Soc. Series C (Appl. Stat.)*, 45(4):516–526, 1996. (ref: p. 95)

[69] K. S. Sarkaria. Tverberg's theorem via number fields. *Israel J. Math.*, 79(2-3):317–320, 1992. (refs: pp. 1, 10)

[70] P. Sarrabezolles. *Colorful Linear Programming.* PhD thesis, 2015. (refs: pp. 21, 35)

[71] P. Sarrabezolles. The colourful simplicial depth conjecture. *J. Combinatorial Theory Ser. A*, 130:119–128, 2015. (ref: p. 7)

[72] A. A. Schäffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM J. Comput.*, 20(1):56–87, 1991. (refs: pp. 25, 60, 61)

[73] A. Schrijver. *Theory of linear and integer programming.* Wiley, 1986. (ref: p. 26)

[74] S.-H. Teng. *Points, Spheres, and Separators: A Unified Geometric Approach to Graph Partitioning.* PhD thesis, Carnegie Mellon University, 1991. (refs: pp. 1, 9, 19)

[75] J. W. Tukey. Mathematics and the picturing of data. In *Proc. Int. Congr. Math.*, volume 2, pages 523–531, 1975. (ref: p. 1)

[76] H. Tverberg. A generalization of Radon's theorem. *J. London Math. Soc*, 41(1):123–128, 1966. (refs: pp. 1, 10)

[77] H. Tverberg. A generalization of Radon's theorem II. *Bull. Australian Math. Soc.*, 24(3):321–325, 1981. (ref: p. 10)

[78] H. Tverberg and S. Vrećica. On generalizations of Radon's theorem and the ham sandwich theorem. *Eur. J. Combin.*, 14(3):259–264, 1993. (ref: p. 10)

[79] L. G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8(2):189–201, 1979. (refs: pp. 71, 98, 103)

[80] U. Wagner. *On k-sets and applications.* PhD thesis, 2003. (ref: p. 14)

[81] R. J. Webster. Another simple proof of Kirchberger's theorem. *J. Math. Anal. and Appl.*, 92(1):299–300, 1983. (ref: p. 12)

[82] G. M. Ziegler. *Lectures on polytopes.* Springer, 1995. (ref: p. 35)

Selbständigkeitserklärung

Ich erkläre hiermit, dass ich alle Hilfsmittel und Hilfen angegeben habe und versichere, auf dieser Grundlage die Arbeit selbständig verfasst zu haben. Die Arbeit habe ich nicht in einem früheren Promotionsverfahren eingereicht.

Berlin, den

Zusammenfassung

Der bunte Satz von Carathéodory ist eine Existenzaussage, die verschiedene Resultate in der konvexen Geometrie nach sich zieht. Dazu gehören unter anderem der Satz von Tverberg, die Existenz von Zentrumspunkten, das erste Selektionslemma und der bunte Satz von Kirchberger. Diese Beweise können als Polynomialzeitreduktionen auf CCP, das zu dem bunten Satz von Carathéodory gehörende algorithmische Problem, interpretiert werden. In dieser Arbeit werden Approximationsalgorithmen und Komplexitätsschranken entwickelt, die aufgrund der Polynomialzeitreduktionen auf CCP auch auf die oben genannten Probleme übertragbar sind.

Seien $C_1, \ldots, C_{d+1} \subset \mathbb{R}^d$ endliche Punktmengen, sodass der Ursprung in der konvexen Hülle jeder Punktmenge C_i, $i \in [d+1]$, enthalten ist. Der bunte Satz von Carathéodory garantiert nun die Existenz einer Auswahl $\boldsymbol{c}_1 \in C_1, \ldots, \boldsymbol{c}_{d+1} \in C_{d+1}$, welche den Ursprung ebenfalls in der konvexen Hülle enthält. CCP beschreibt dann das Problem, diese Auswahl zu berechnen. Da immer eine Lösung existiert und eine Kandidatenlösung in Polynomialzeit überprüfbar ist, liegt CCP in der Komplexitätsklasse *totale Funktionen NP* (TFNP), die Klasse der NP-Suchprobleme für die immer eine Lösung existiert. In dieser Arbeit wird gezeigt, dass CCP im Schnitt zweier wichtiger Unterklassen von TFNP enthalten ist: der Komplexitätsklasse *Polynomialzeitparitätsargument in gerichteten Graphen* (PPAD) und in der Komplexitätsklasse *polynomielle lokale Suche* (PLS). Die Formulierung von CCP als PPAD-Problem basiert auf einem neuen konstruktiven Beweis des bunten Satzes von Carathéodory durch Sperners Lemma. Des Weiteren wird gezeigt, dass schon eine kleine Änderung in der Definition von CCP zu einem PLS-vollständigen Problem führt.

Im zweiten Teil der Arbeit werden verschiedene konstruktive Resultate vorgestellt. Zuerst wird das Approximationsproblem betrachtet, indem mehr als ein Punkt von jeder Menge C_i ausgewählt werden kann. Dies ist mit den Polynomialzeitreduktionen auf CCP kompatibel. Es wird gezeigt, dass für jedes feste $\varepsilon > 0$ eine Auswahl C mit maximal $\lceil \varepsilon d \rceil$ Punkten von jeder Menge C_i in Polynomialzeit gefunden werden kann, sodass $\mathbf{0} \in \operatorname{conv}(C)$. Zusätzlich wird ein verwandtes Approximationsproblem betrachtet, in dem die Eingabe aus $k < d+1$ Mengen C_i besteht mit $\mathbf{0} \in \operatorname{conv}(C_i)$ für $i \in [k]$ und eine Auswahl C mit $\mathbf{0} \in \operatorname{conv}(C)$ gefunden werden soll, die maximal $\lceil (d+1)/k \rceil$ Punkte von jeder Menge C_i enthält. Die Existenz von C ist eine direkte Implikation des bunten Satzes von Carathéodory. Mithilfe von linearen Programmen ist es möglich, dieses Problem für den Fall $k = 2$ in schwach polynomieller Zeit zu lösen. Des Weiteren wird gezeigt, dass CCP exakt in quasi-polynomieller Zeit gelöst werden kann, falls die Eingabe aus $\operatorname{poly}(d)$ Mengen besteht, anstatt nur aus $d+1$. Abschließend wird das Problem der simplizialen Tiefe betrachtet. Die *simpliziale Tiefe* $\sigma_P(\boldsymbol{q})$ eines Punktes $\boldsymbol{q} \in \mathbb{R}^d$ bezüglich einer Menge $P \subset \mathbb{R}^d$ ist die Anzahl aller verschiedener d-Simplexe mit Ecken in P, die $\boldsymbol{q}$ enthalten. Es wird gezeigt, dass $\sigma_P(\boldsymbol{q})$ in $\widetilde{O}\left(n^{d/2+1}\right)$ Zeit mit hoher Wahrscheinlichkeit $(1+\varepsilon)$-approximiert werden kann für $\varepsilon > 0$ fest. Für den Fall, dass die Dimension Teil der Eingabe ist, wird bewiesen, dass das Problem #P-vollständig und W[1]-schwer ist.