Tropical Bisectors and Voronoi Diagrams

Francisco Criado¹ · Michael Joswig^{1,2} · Francisco Santos³

Received: 3 July 2019 / Revised: 30 January 2021 / Accepted: 27 May 2021 /

Published online: 7 September 2021

© The Author(s) 2021

Abstract

In this paper we initiate the study of tropical Voronoi diagrams. We start out with investigating bisectors of finitely many points with respect to arbitrary polyhedral norms. For this more general scenario we show that bisectors of three points are homeomorphic to a non-empty open subset of Euclidean space, provided that certain degenerate cases are excluded. Specializing our results to tropical bisectors then yields structural results and algorithms for tropical Voronoi diagrams.

Keywords Polyhedral norms · Polytropes · Randomized algorithms · Voronoi diagrams

Mathematics Subject Classification 14T15 · 46B20 · 52B55

Communicated by Evelynne Hubert.

F. Criado has been supported by Berlin Mathematical School and Einstein Foundation Berlin (EVF-2015-230). M. Joswig has been supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - The Berlin Mathematics Research Center MATH⁺ (EXC-2046/1, Project ID 390685689); "Symbolic Tools in Mathematics and their Application" (TRR 195/2, Project-ID 286237555); "Facets of Complexity" (GRK 2434). F. Santos has been supported by the Einstein Foundation Berlin (EVF-2015-230) and by Grants MTM2017-83750-P/AEI/10.13039/501100011033 and PID2019-106188GB-I00/AEI/10.13039/501100011033 of the Spanish State Research Agency.

✉ Michael Joswig
joswig@math.tu-berlin.de

Francisco Criado
criado@math.tu-berlin.de

Francisco Santos
francisco.santos@unican.es

¹ TU Berlin, Chair of Discrete Mathematics/Geometry, Berlin, Germany

² Chair of Discrete Mathematics/Geometry; Max-Planck Institute for Mathematics in the Sciences, Leipzig, Germany

³ Departamento de Matemáticas, Estadística y Computación, Univ. de Cantabria, Santander, Spain

1 Introduction

One early route to the success of tropical geometry is based on the tropicalization of classical algebraic varieties defined over some valued field. Key examples include Mikhalkin's correspondence principle, which relates tropical plane curves with classical complex algebraic curves [24], or the tropical Grassmannians of Speyer and Sturmfels [25]. In all of this the focus lies on the combinatorial properties of tropical varieties, which are ordinary polyhedral complexes.

More recently, however, tropical semi-algebraic sets and their intrinsic geometry came into the picture; see [1, 15]. For instance, their metric properties appear in [2] as a tool to show that standard versions of the interior point method of linear programming exhibit an exponential complexity in the unit cost model. The proof of that result is based on translating metric data on a family of tropical linear programs into curvature information about the central paths of their associated ordinary linear programs. Similarly, tropical analogs of isoperimetric (or isodiametric) inequalities have been studied in [8], where a tropical volume is defined that corresponds to an “energy gap” in mathematical physics [19]. Another example is the statistical analysis of phylogenetic trees by Lin, Monod and Yoshida [21].

We feel that all this calls for a more systematic investigation of metric properties of tropical varieties. Starting from first principles, this naturally leads to tropical Voronoi diagrams. The *tropical distance* between two points $a, b \in \mathbb{R}^{d+1}$ is

$$\text{dist}(a, b) = \max_{i \in [d+1]} (a_i - b_i) - \min_{j \in [d+1]} (a_j - b_j) = \max_{i, j \in [d+1]} (a_i - b_i - a_j + b_j) . \quad (1)$$

It does not depend on choosing min or max as the tropical addition. The map $\text{dist} : \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ is non-negative, symmetric, and it satisfies the triangle inequality. Moreover, it is homogeneous, so it induces a norm on the *tropical d -torus* $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1} \cong \mathbb{R}^d$, where $\mathbf{1} = (1, \dots, 1)$ denotes the all ones vector. The *tropical Voronoi region* of a site $s \in S$ with respect to a set S comprises those points in $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$ to which s is the nearest among all sites in S , with respect to dist . The *tropical Voronoi diagram* $\text{Vor}(S)$ is the cell decomposition of $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$ into Voronoi regions. Tropical Voronoi diagrams are a special case of Voronoi diagrams for polyhedral norms, a classical topic in convexity and computational geometry; cf. [4, Sect. 7.2] or [23, Sect. 4].

The intersection of two or more Voronoi regions is part of a *bisector*, i.e., the locus of points which are equidistant to a given set. For instance, in the Euclidean case the bisector of two points is a degenerate quadric which agrees with an affine hyperplane as a set. In the tropical setting, the bisector of two points can also be described as part of a tropical hypersurface, but this is now of degree $d + 1$; see Proposition 4. Further, in the tropical case two points may already produce degenerate bisectors (which may contain, e.g., full-dimensional pieces), whereas the first degenerate case in the Euclidean metric arises for three points. So tropical Voronoi diagrams behave quite differently from Euclidean Voronoi diagrams.

Yet there are also similarities. A key structural result is that the tropical Voronoi regions are star convex and can be described as unions of finitely many ordinary poly-

hedra; see Proposition 1 and Theorem 6. We prove a second main result, Theorem 3, for the more general case of an arbitrary polyhedral norm in $\mathbb{R}^d$: the bisector of any three points in weak general position is homeomorphic to an open subset of $\mathbb{R}^{d-2}$. Our proof generalizes the arguments from [13], [14], where a similar result was proved for smooth norms in $d = 2, 3$. However, the global topology of tropical bisectors of three or more points can be radically different from the topology of the classical bisectors. For instance, tropical bisectors are sometimes disconnected and, more strongly, $d + 1$ points can have more than one circumcenter. This may happen even in general position; see Examples 1 and 3. We do not know if bisectors may have nontrivial higher Betti numbers, but we suspect they can; see Theorem 4.

Another contribution in our paper is a randomized incremental algorithm for computing the tropical Voronoi diagram of n points in general position in $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$ with an expected running time of order $O(n^d \log n)$, for fixed dimension d ; see Theorem 10. Euclidean Voronoi diagrams of finite point sets can be explained fully in terms of ordinary convex polyhedra and convex hull algorithms; see [4,7]. We do not know if there is a tropical analog.

Amini and Manjunath [3] study the Voronoi diagram of a lattice with respect to the following asymmetric version of the tropical distance:

$$\text{dist}(a, b) = \max_{i \in [d+1]} (a_i - b_i).$$

As they show in [3, Lemma 4.7], this is the polyhedral distance obtained taking as unit ball the standard simplex. Their motivation comes from work of Baker and Norine [5] on a Riemann–Roch theorem for graphs, which implies a Riemann–Roch theorem for tropical curves.

Our paper is organized as follows. The short Sect. 2, in which we verify that the tropical distance is induced by a polyhedral norm and discuss the combinatorics of the tropical unit ball, sets the stage. In Sect. 3 we collect our general structural results on bisectors and Voronoi diagrams. The results in this section are proved for general polyhedral norms, but all our examples address the tropical case. A subtle point is the right concept of “general position”. In fact, we distinguish between *weak general position* which prevents bisectors to contain full-dimensional parts (see Proposition 2), and a stronger *general position* which is defined via stability of facets in the bisector under perturbation of the sites. For instance, the bisector of any number k of points in general position in $\mathbb{R}^d$ is either empty or a polyhedral complex of *pure* dimension $d + 1 - k$; see Corollary 2. As a special case, the bisector of $d + 1$ points in $\mathbb{R}^d$ in general position is finite. Section 4 returns to the tropical case. We specialize our results on bisectors in general polyhedral norms, and we show that the combinatorial types of tropical bisectors of two points are classified in terms of a certain polyhedral fan related to the tropical unit ball and the braid arrangement; see Theorems 5 and 6. This is related to work of Develin [9] on the moduli of tropically collinear points. Finally, in Sect. 5 we discuss algorithms. This includes a tropical variant of Fortune’s beach line algorithm [10] for planar Voronoi diagrams as well as the aforementioned algorithm in arbitrary dimension.

2 The Tropical Unit Ball

The unit ball with respect to the tropical distance function defined in (1) is

$$\begin{aligned}\mathbb{B}^d &= \left\{x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \mid \text{dist}(x, 0) = 1\right\} = \bigcap_{i \neq j} \left\{x \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \mid x_i - x_j \leq 1\right\} \\ &= \frac{1}{2} \text{conv} \left(\{\pm 1\}^{d+1} \setminus \{\pm \mathbb{1}\} \right) + \mathbb{R}\mathbb{1} .\end{aligned}\quad (2)$$

In this way, $\mathbb{B}^d$ is a polytope in the tropical torus $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. We also write $\mathbb{B}^d(a, r)$ for the tropical ball with center a and radius r . All tropical balls result from scaling and translating $\mathbb{B}^d$. In fact, the tropical norm agrees with polyhedral norm with respect to the tropical unit ball, in the sense of Sect. 3. Such distances are called *convex distance functions* in [4, Sect. 7.2]; see also [12–14].

Both the inequality and the vertex descriptions of $\mathbb{B}^d$ in Eq. (2) are non-redundant:

- $\mathbb{B}^d$ has $d(d+1)$ facets. Each facet corresponds to a choice of coordinates achieving the maximum and the minimum.
- $\mathbb{B}^d$ has $2^{d+1} - 2$ vertices. Each vertex corresponds to a (nontrivial) partition of the coordinates into maxima and minima. For example, $\mathbb{B}^2$ is a hexagon and $\mathbb{B}^3$ is a rhombic dodecahedron.

The vertex description also shows that $\mathbb{B}^d$ equals the projection of the $(d+1)$ -dimensional regular cube $[-1, 1]^{d+1}$ in $\mathbb{R}^{d+1}$ along the direction $\mathbb{1}$. That is, $\mathbb{B}^d$ is a zonotope with $d+1$ generators in general position, and all its faces are parallelepipeds. These generators correspond to the $d+1$ coordinate directions in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. This suggests a combinatorial way to specify the faces of $\mathbb{B}^d$: Each face F can be written as a Minkowski sum

$$F = \sum_{i=1}^{d+1} s_i ,$$

where each s_i is one of $\{-e_i\}$, $[-e_i, e_i]$ or $\{+e_i\}$. We say that F is of *type* (F_-, F_*, F_+) if

$$\begin{aligned}F_- &= \{i \in [d+1] : s_i = \{-e_i\}\} , \\ F_* &= \{i \in [d+1] : s_i = [-e_i, e_i]\} , \\ F_+ &= \{i \in [d+1] : s_i = \{e_i\}\} .\end{aligned}\quad (3)$$

Conversely, a partition of $[d+1]$ into three parts F_-, F_*, F_+ corresponds to a face of $\mathbb{B}^d$ if and only if neither F_- nor F_+ is empty. Moreover, the dimension of F equals the cardinality of F_* . In particular, the vertices of $\mathbb{B}^d$ correspond to the $2^{d+1} - 2$ ways of partitioning $[d+1]$ into two non-empty subsets. The facets of $\mathbb{B}^d$ correspond to the $d(d+1)$ ways of choosing an ordered pair from $[d+1]$, without repetition (Fig. 1).

Remark 1 The zonotope $\mathbb{B}^d$ is dual to an arrangement of $d+1$ linear hyperplanes in general position in $\mathbb{R}^d$, oriented so that the intersection of all positive half-spaces is

Fig. 1 The tropical 3-ball $\mathbb{B}^3$, with the conical hull of a facet highlighted

empty. In particular, its face lattice is the same as the lattice of nonzero covectors of the unique totally cyclic oriented matroid of rank d with $d + 1$ elements. Covectors of an oriented matroid are usually written as (V_-, V_0, V_+) but in our context we prefer to use $*$ instead of zero meaning that the corresponding coordinate is not fixed.

Remark 2 Another general description of $\mathbb{B}^d$ is that it equals the (ordinary) Voronoi cell of the lattice of type A_d (i.e., the triangular lattice for $d = 2$ and the face centered cubic lattice (FCC) for $d = 3$).

Similarly, $\mathbb{B}^d$ is the polytope polar to the difference body $T - T$ of a regular d -simplex T . This description shows that $\mathbb{B}^d$ is the same as the polytope U_d that appears in Makeev’s conjecture. See, e.g., [26, Conjecture 21.3.2].

3 Bisectors in Polyhedral Norms

Throughout this section we work in the general framework of Minkowski norms; see [4, Sect. 7.2], [12], [23]. Consider a polytope $K \subset \mathbb{R}^d$ with the origin in its interior. Let $\text{dist}(a, b)$ be the unique scaling factor $\alpha > 0$ such that $b - a \in \partial(\alpha K)$. Then, dist satisfies the triangle inequality, is invariant under translation, and homogeneous under scaling. If $K = -K$ then $\text{dist}(a, b) = \text{dist}(b, a)$ and $\text{dist}(0, \cdot)$ is a norm in $\mathbb{R}^d$ in the usual sense. We allow $K \neq -K$, whence $\text{dist}(a, b) \neq \text{dist}(b, a)$ in general, but we still call it a norm. Bisectors and Voronoi diagrams for these norms have been studied in computational geometry [4, Sect. 7.2], [23, Sect. 4].

For any finite point set S we define its *bisector*:

$$\text{bis}(S) := \left\{ x \in \mathbb{R}^d \mid \text{dist}(a, x) = \text{dist}(b, x) \text{ for } a, b \in S \right\}.$$

Following the computational geometry tradition we will often call the elements of S the *sites*. Although some of our results also hold for general convex bodies, for simplicity we assume K to be a polytope. We denote by $\mathcal{F}(K)$ the face fan of K . The

Fig. 2 Left: A point, p , in the tropical bisector of a and b . Right: The analogous picture, in classical geometry

norm $\text{dist}(0, \cdot)$ is linear in each of these cones, so we write

$$\text{bis}_{(F_1, \dots, F_k)}(\{a_1, \dots, a_k\}) = \text{bis}(\{a_1, \dots, a_k\}) \cap a_1 + F_1 \cap \dots \cap a_k + F_k,$$

for the intersection of the bisector with a choice of cones $F_i \in \mathcal{F}(K)$. Each cell of the form $\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k)$ is the intersection of the polyhedron $(a_1 + F_1) \cap \dots \cap (a_k + F_k)$ with an affine subspace, which implies it is itself a polyhedron. As a consequence:

Proposition 1 *Let K be a polytope with the origin in its interior, and let dist be the corresponding Minkowski norm. Let $S = \{a_1, \dots, a_k\} \subset \mathbb{R}^d$ be a finite point set. Then the bisector $\text{bis}(\{a_1, \dots, a_k\})$ is a polyhedral complex whose cells are the polyhedra*

$$\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k)$$

for all choices of $F_1, \dots, F_k \in \mathcal{F}(K)$.

Proof The family of polyhedra

$$\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k), \text{ with } F_1, \dots, F_k \in \mathcal{F}(K),$$

forms a polyhedral complex since

$$\text{bis}_{(F_1, \dots, F_k)}(S) \cap \text{bis}_{(F'_1, \dots, F'_k)}(S) = \text{bis}_{(F_1 \cap F'_1, \dots, F_k \cap F'_k)}(S).$$

That polyhedral complex covers the entire bisector since for each point $p \in \text{bis}(S)$ and for each i , the point a_i must lie in some face F_i of $p - \text{dist}(a_i, p)K$. $\square$

Our primary example is the case where $K = \mathbb{B}^d$ is the tropical ball. In Fig. 2 the point p , which is generic within the bisector of a and b , lies in the facet $\text{bis}_{(-*, +)}(a, b)$.

For the purpose of drawing pictures, notice that any three vectors $v_1, v_2, v_3 \in \mathbb{R}^2$ with $v_1 + v_2 + v_3 = 0$ define a map from $\mathbb{R}^3/\mathbb{R}\mathbf{1}$ to $\mathbb{R}^2$ via $e_i \mapsto v_i$. While $v_1 = (1, 0)$, $v_2 = (0, 1)$, $v_3 = (-1, -1)$ is a common base for diagrams in tropical geometry, for our pictures we settle for the more symmetric *isometric view* where our base is:

$$v_1 = \left(-\sin \frac{2\pi}{3}, \cos \frac{2\pi}{3} \right), \quad v_2 = \left(\sin \frac{2\pi}{3}, \cos \frac{2\pi}{3} \right), \quad v_3 = (0, 1) .$$

Remark 3 If $\text{bis}(a_1, \dots, a_k) = \emptyset$ then $\text{bis}(a'_1, \dots, a'_k) = \emptyset$ for every choice of $a'_1, \dots, a'_k$ sufficiently close to $a_1, \dots, a_k$, by continuity of $\text{dist}(\cdot, \cdot)$.

3.1 Weak General Position and General Position

Definition 1 (General position) A finite point set $S \subset \mathbb{R}^d$ is in *general position* with respect to K , if for every subset $\{a_1, \dots, a_k\} \subset S$ there are neighborhoods U_i of each a_i such that for every choice of $\{a'_1, \dots, a'_k\}$ with $a'_i \in U_i$ and for every choice of maximal cones $F_1, \dots, F_k \in \mathcal{F}(K)$ we have

$$\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k) = \emptyset \iff \text{bis}_{(F_1, \dots, F_k)}(a'_1, \dots, a'_k) = \emptyset .$$

Moreover, the set S is in *weak general position* if no pair of points $a, b \in S$ lies in a hyperplane parallel to a facet of K .

Remark 4 As the name suggests, “weak general position” is implied by “general position” (see Corollary 1). A yet stronger notion would arise requiring stability not only of facets but also of lower-dimensional cells in bisectors; that is, allowing lower-dimensional cones from $\mathcal{F}(K)$ for the F_i in the definition of general position. But this intermediate notion of “general position” is the most appropriate for our purposes and the algorithms in Sect. 5. As a first indication, Theorem 1 provides a local characterization.

By Proposition 1 bisectors are polyhedral complexes, and thus each cell has a dimension.

Proposition 2 Two points a, b are in weak general position if and only if $\text{bis}(a, b)$ does not contain full-dimensional cells.

Proof Since $\text{bis}_{(F, F')}(a, b) \subset (a + F) \cap (b + F')$, for it to be d -dimensional we need F and F' to be cones of facets. We also need $F = F'$, so that $\text{dist}(a, \cdot)$ and $\text{dist}(b, \cdot)$ have the same gradient on $(a + F) \cap (b + F)$, and we need $b - a$ to be parallel to the facet, so that $\text{dist}(a, \cdot) = \text{dist}(b, \cdot)$ on $(a + F) \cap (b + F)$.

Conversely, if $b - a$ is parallel to a facet of K with cone F then

$$\text{bis}_{(F, F)}(a, b) = (a + F) \cap (b + F) ,$$

and this is d -dimensional. □

Corollary 1 *General position implies weak general position.*

Proof Suppose $a - b$ is parallel to a facet of K , so that $\text{bis}_{(F,F)}(a, b)$ is full-dimensional, where F is the cone of that facet. Taking b' close to b but away from the hyperplane parallel to the facet makes $\text{bis}_{(F,F)}(a, b')$ empty. Now the claim follows from Proposition 2. $\square$

Theorem 1 *Let $S = \{a_1, \dots, a_k\} \subseteq \mathbb{R}^d$ and for each a_i choose a maximal cone $F_i \in \mathcal{F}(K)$. Let $Q := (a_1 + F_1) \cap \dots \cap (a_k + F_k)$, let $\lambda_{F_i}(x)$ be the linear function that restricts to $\text{dist}(0, x)$ on F_i , and let H be the affine subspace defined by $\lambda_{F_1}(x - a_1) = \dots = \lambda_{F_k}(x - a_k)$.*

Then, the following conditions are equivalent:

1. *There are neighborhoods U_i of each a_i such that for any choice of $a'_i \in U_i$ the polyhedron $\text{bis}_{(F_1, \dots, F_k)}(a'_1, \dots, a'_k)$ is not empty.*
2. (a) *Q is full-dimensional and H intersects its interior; and*
 (b) *the $k - 1$ functions $\lambda_{F_i} - \lambda_{F_1}$ for $i = 2, \dots, k$ are linearly independent.*

Since

$$\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k) = Q \cap H \quad (4)$$

condition (a) is equivalent to “ $\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k)$ meets the interior of Q ”.

Proof For the implication from “1” to “2”, let us first show that “1” forces Q to be full-dimensional. Aiming for a contradiction, we assume that Q is contained in the boundary of one of the cones $a_i + F_i$. That is, the polyhedron $Q_i := \bigcap_{j \neq i} (a_j + F_j)$ does not meet the interior of $a_i + F_i$, for some i . Then any a'_i in the interior of $a_i + F_i$ yields $(a'_i + F_i) \cap Q_i = \emptyset$. Hence $\text{bis}_{(F_1, \dots, F_k)}(a'_1, \dots, a'_k) = \emptyset$, where $a'_j = a_j$ for $j \neq i$. This contradicts “1” and shows that Q is full-dimensional.

To see that $\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k)$ must intersect the interior of Q , suppose we are given neighborhoods U_i as in “1”; recall (4). For each i , choose v_i in the interior of F_i and such that $\lambda_i(v_i) = 1$. Let $a'_i := a_i + \epsilon v_i$, where $\epsilon > 0$ is taken small enough so that $a'_i \in U_i$. Observe that our choice of v_i makes the affine subspace defined by $\lambda_{F_1}(x - a'_1) = \dots = \lambda_{F_k}(x - a'_k)$ agree with H . Let $Q' := (a'_1 + F_1) \cap \dots \cap (a'_k + F_k)$, which lies in the interior of Q . We have

$$\text{bis}_{(F_1, \dots, F_{d+1})}(a'_1, \dots, a'_{d+1}) = Q' \cap H.$$

By “1” this is not empty, and thus H intersects the interior of Q .

For condition (b) observe that H can equivalently be defined by the $k - 1$ affine equalities

$$(\lambda_{F_i} - \lambda_{F_1})(x) = \lambda_{F_i}(a_i) - \lambda_{F_1}(a_1) \quad \text{for } i = 2, \dots, k. \quad (5)$$

If the left-hand sides are linearly dependent, then one of the $k - 1$ equations, say the i th one, is redundant. But then choosing a point $a'_i \in U_i$ with $\lambda_{F_i}(a'_i) \neq \lambda_{F_i}(a_i)$

(and letting $a'_j = a_j$ for the rest) renders the system of equations infeasible. Hence $\text{bis}_{(F_1, \dots, F_{d+1})}(a'_1, \dots, a'_{d+1}) = \emptyset$, contradicting “1”.

We now show that “2” implies “1”. Consider arbitrary points a'_i , and let $Q'_i = \bigcap_i (a'_i + F_i)$. Further let H' be the affine subspace defined by

$$\lambda_{F_1}(x - a'_1) = \dots = \lambda_{F_k}(x - a'_k) ,$$

so that

$$\text{bis}_{(F_1, \dots, F_k)}(a'_1, \dots, a'_k) = Q' \cap H' .$$

We want to show that if each a'_i is sufficiently close to the corresponding a_i for all i then $Q' \cap H'$ is not empty.

Condition (b) says that H' is $(d+1-k)$ -dimensional (and parallel to H) for every choice of a'_i s and that it varies continuously with the choice. Condition (a) says that Q' stays full-dimensional if the a'_i s are sufficiently close to the original a_i s, and that it also varies continuously with the choice, in the following strong sense: consider a description of each $a_i + F_i$ by a finite system of linear inequalities. Then $a'_i + F_i$ is defined by a system with the same linear functions and with right-hand sides varying continuously with the a'_i s.

Thus, if each a'_i is close to a_i , then Q' and H' are a full-dimensional polyhedron and a $(d+1-k)$ -dimensional affine subspace close to Q and H respectively. Since, by (a), H intersects the interior of Q , continuity implies that H' still intersects the interior of Q' when a'_i is close enough to a_i for each i . In particular, $\text{bis}_{(F_1, \dots, F_k)}(a'_1, \dots, a'_k)$ is not empty. $\square$

Corollary 2 *The bisector of k points in general position is either empty or pure of dimension $d + 1 - k$. In particular, the bisector of $d + 1$ points in general position is finite, and this is empty for more than $d + 1$ points.*

Proof Every maximal non-empty cell $\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k)$ is the intersection of the polyhedron Q with the affine subspace H of Theorem 1. That result implies that H has dimension $d + 1 - k$ and meets the interior of the full-dimensional polyhedron Q . Thus, the cell has dimension $d + 1 - k$. $\square$

Corollary 3 *If every subset of at most $d + 2$ points in S is in general position then so is S .*

Corollary 4 *For any $n \geq 1$, the sequences of n points in $\mathbb{R}^d$ in general position form an open dense subset of $(\mathbb{R}^d)^n$.*

Proof Let $S \subset \mathbb{R}^d$ be a set of cardinality n . Then S is in general position if and only if, for each subset $\{a_1, \dots, a_k\} \subset S$ and maximal cones $F_1, \dots, F_k \in \mathcal{F}(K)$, the polyhedron $\text{bis}_{(F_1, \dots, F_k)}(a_1, \dots, a_k)$ either is empty or satisfies condition (1) of Theorem 1. Since there are finitely many choices of $\{a_1, \dots, a_k\}$ and $\{F_1, \dots, F_k\}$, it suffices to prove the statement for one such choice.

For openness: Emptiness is an open condition because of the continuity of $\text{dist}(\cdot, \cdot)$. On the other hand, Theorem 1 says that condition (1) is equivalent to (2), which is open

by the arguments in the proof (namely, the fact that both Q and H depend continuously on the sites).

For density, we are going to show that if $Q \cap H$ is not empty but fails to satisfy condition (a) or (b) then the set $\{a_1, \dots, a_k\}$ lies in one of finitely many linear hyperplanes in $(\mathbb{R}^d)^k$. To emphasize that Q and H depend on the choice of sites we denote them $Q(a_1, \dots, a_k)$ and $H(a_1, \dots, a_k)$.

If (b) fails for $\{a_1, \dots, a_k\}$ then the linear system (5) defining $H(a_1, \dots, a_k)$ is feasible but overdetermined, which implies a linear relation, depending solely on $F_1, \dots, F_k$, among the right-hand sides $\lambda_{F_i}(a_i) - \lambda_{F_1}(a_1)$. The relation is not tautological on the a_i s since, as shown in the Proof of Theorem 1, it is easy to construct a point set $\{a'_1, \dots, a'_k\}$ with $H(a'_1, \dots, a'_k) = \emptyset$.

For (a), consider the inequality descriptions of the cones F_i and translate them to obtain an inequality description of $Q(a_1, \dots, a_k)$ as the feasibility region of a system $\mathcal{S}(a_1, \dots, a_k)$ of affine inequalities with fixed gradients and with right-hand sides parameterized linearly by the a_i s. If (a) fails for $\{a_1, \dots, a_k\}$ then one of two things happen:

- $Q(a_1, \dots, a_k)$ is non-empty but not full-dimensional. Consider a minimal subsystem of $\mathcal{S}(a_1, \dots, a_k)$ that already defines a non-full-dimensional feasibility region. Minimality implies that this feasibility region is an affine subspace of $\mathbb{R}^d$ and that turning the inequalities to equalities produces an over-determined subsystem. This implies, as in the previous case, a linear relation among the a_i s.
- $Q(a_1, \dots, a_k)$ is full-dimensional but $Q(a_1, \dots, a_k) \cap H(a_1, \dots, a_k)$ is contained in its boundary. This implies that $Q(a_1, \dots, a_k) \cap H(a_1, \dots, a_k)$ is contained in a facet of $Q(a_1, \dots, a_k)$. Let $H_0(a_i)$ be the hyperplane containing that facet. Note that the hyperplane H_0 depends only on one a_i since it comes from the description of one of the cones $a_i + F_i$. Then, adding to the the linear system (5) the equation defining $H_0(a_i)$ produces again an over-determined system, hence there is a linear relation among the a_i s.

The relations for condition (a) are not tautological on the a_i s since we can easily make $Q(a'_1, \dots, a'_k)$ full-dimensional and $H(a'_1, \dots, a'_k)$ intersect its interior as follows: choose each a'_i in the interior of $-F_i$ and such that $\lambda_{F_i}(a_i) = -1$, so that both H and the interior of Q contain the origin. The relations are finitely many since we have at most one for each subsystem (respectively equation) of the system $\mathcal{S}(a_1, \dots, a_k)$, which depends only on the choice of F_i s. $\square$

For the case of the tropical norm, condition (b) admits a nice combinatorial characterization. Observe that a choice of facets $F_1, \dots, F_k \in \mathcal{F}(\mathbb{B}^d)$ can be encoded as a directed graph on the vertex set $[d + 1]$ and with an arc a_i going from the coordinate that is minimized at F_i to the coordinate that is maximized at F_i , for $i = 1, \dots, k$. We denote this graph $G(F_1, \dots, F_k)$.

Proposition 3 *For the case of the tropical norm, condition 2.(b) of Theorem 1 holds if and only if the graph $G(F_1, \dots, F_k)$ either has no (undirected) cycle or it has a unique cycle and it is unbalanced; that is, the number of arcs in one direction is different from the other direction.*

Proof A cycle in $G(F_1, \dots, F_k)$ is equivalent to a linear dependence among the corresponding linear functions λ_{F_i} s, by simply adding them with signs corresponding to the direction of the arcs along the cycle. If the cycle is balanced then λ_{F_1} can be subtracted from each λ_{F_i} so that the corresponding functions $(\lambda_{F_i} - \lambda_{F_1})$ are also dependent. The same thing can be done if $G(F_1, \dots, F_k)$ has two different (unbalanced) cycles, since a linear combination of the two corresponding dependences can be made balanced.

Conversely, any linear dependence among the functions $(\lambda_{F_i} - \lambda_{F_1})$ corresponds to a balanced dependence among the corresponding λ_{F_i} . The latter either corresponds to a balanced circuit in the graph or decomposes into two (or more) linear dependences with distinct supports. $\square$

As a consequence of Corollary 2 the bisector of a set S of $d + 1$ points in general position is a finite set of points, which we call *circumcenters* of S . In dimension two, three points in (weak) general position have a most one circumcenter, as we show in Corollary 6 below. In higher dimension the same is known to be false for other polytopal norms [13], and here is an example for the tropical norm:

Example 1 (Non-uniqueness of circumcenters) Let us consider the four points $a_1 = (0, 2, 3, 3)$, $a_2 = (0, 4, 2, 2)$, $a_3 = (2, 4, 1, 1)$ and $a_4 = (4, 0, 2, 2)$. Their bisector contains the points $x = (0, 0, 1, -1)$ and $y = (0, 0, -1, 1)$. Indeed, both x and y are at distance 4 from all the a_i 's since we have

$$\begin{aligned} a_1 - x &= (0, 2, 2, 4), & a_1 - y &= (0, 2, 4, 2), \\ a_2 - x &= (0, 4, 1, 3), & a_2 - y &= (0, 4, 3, 1), \\ a_3 - x &= (2, 4, 0, 2), & a_3 - y &= (2, 4, 2, 0), \\ a_4 - x &= (4, 0, 1, 3), & a_4 - y &= (4, 0, 3, 1). \end{aligned}$$

The points $a_1, \dots, a_4$ are not in weak general position, as they lie in the plane $x_3 - x_4 = 0$. However, they satisfy conditions (a) and (b) of Theorem 1 for the polytopes Q_x and Q_y containing the circumcenters x and y : For condition (b) observe that the digraphs corresponding to x and y are, respectively, $\{14, 12, 32, 21\}$ and $\{13, 12, 42, 21\}$. They both have a single cycle, $\{12, 21\}$, which is unbalanced. Condition (a) follows from the fact that x (resp. y) is in the interior of all the cones whose intersection defines Q_x (resp. Q_y). This is equivalent to the fact that all the vectors $a_i - x$ and $a_i - y$ have a unique maximum and a unique minimum entries.

Since condition (b) does not depend on the sites and condition (a) is, by Theorem 1, open, any perturbation of the sites will still produce at least two circumcenters. In particular, by Corollary 4, there are sites in general position for which this happens.

3.2 Halfspheres, Sectors, and the Bisector of Two Points

The topology of a bisector is closely related to the following partition of ∂K . Let $S \subseteq \mathbb{R}^d$ be a finite set of sites, and pick $a, b \in S$ distinct. The *open halfsphere in the direction of $b - a$* , denoted as $H(b - a)$, is the set of points in ∂K whose exterior normal cone is contained in $(b - a)^\vee := \{\lambda \mid \lambda(b - a) > 0\}$. Informally, $H(b - a)$ are

the faces of K that “can be seen” from the direction $(b - a)$. For a fixed site $a \in S$, the *sector of a* is the set

$$H_S(a) = \bigcap_{b \in S \setminus \{a\}} H(b - a) .$$

We denote $\mathcal{H}_S := \{H_S(a) \mid a \in S\}$. Observe that $H(b - a)$ and, hence, $H_S(a)$, are open in ∂K .

Lemma 1 *Let $F_1, \dots, F_m$ be the facets of K and let $\lambda_{F_i}(x) \leq 1$ be the valid linear inequality defining F_i . Then, for each $a \in S$,*

$$H_S(a) = \operatorname{relint} \left(\bigcup \{F_i \mid \lambda_{F_i}(a) < \lambda_{F_i}(b) \text{ for all } b \in S \setminus \{a\}\} \right) .$$

In particular, $H_S(a) \cap H_S(b) = \emptyset$ for every $a, b \in S$ and, if S is in weak general position,

$$\bigcup_{a \in S} \overline{H_S(a)} = \partial K .$$

where $\overline{H_S(a)}$ denotes the topological closure of $H_S(a)$.

Proof It is clear from the definition that $H_S(a)$ contains the relative interior of every F_i with $\lambda_{F_i}(a) < \lambda_{F_i}(b)$ for $b \in S \setminus a$. By convexity of the cones $\{\lambda \mid \lambda(b - a) > 0\}$ for each a, b , $H(b - a)$ (hence $H_S(a)$) also contains the relative interior of every lower dimensional face contained only in such facets. This proves the first formula. The second part follows from the first and the fact that in weak general position the minimum of each λ_{F_i} is attained at a single point of S . $\square$

Remark 5 Assuming weak general position, Lemma 1 allows us to think of $\mathcal{H}_S$ as a labeling of the facets of K by the elements of S or, equivalently, as a map $\mathcal{F}(K) \rightarrow S$. If K is centrally symmetric, then each pair of opposite facets F and $-F$ belong one to $H(b - a)$ and the other to $H(a - b)$. If K is not, we can still guarantee that $H(a - b)$ is never empty, and always disjoint from $H(b - a)$. As a consequence, $H(b - a)$ (and hence $H_S(a)$) cannot contain all the (relative interiors of) facets of K .

For the case $K = \mathbb{B}^d$ of the tropical ball this partition of the facets translates into something more meaningful. Recall (see Proposition 3 and the paragraph before it) that facets of $\mathbb{B}^d$ can be represented as the arcs in the complete digraph on $d + 1$ nodes. In particular, $\mathcal{H}_S$ colors the arcs by the points of S . Then:

1. Each coloring is a partial ordering of vertices, i.e. there is no monochromatic cycle.
2. For the case of two points in general position, the two colors are opposite acyclic tournaments. In particular, there is a bijection between the possible halvespheres $H(b - a)$ and the total orderings of $d + 1$ elements.

Theorem 2 ([13] for $d = 3$) *Let $a, b \in \mathbb{R}^d$ be in weak general position. Then the central projection from a induces a homeomorphism between $\operatorname{bis}(a, b)$ and $a + H(b - a)$. Hence, $\operatorname{bis}(a, b)$ is homeomorphic to $\mathbb{R}^{d-1}$.*

Proof Let us first show that $\text{bis}(a, b)$ is contained in $a + \text{cone}(H(b - a))$. To seek a contradiction, let $c \in \text{bis}(a, b)$ such that $c - a \notin \text{cone}(H(b - a))$. This implies that the smallest ball centered at a that contains c touches it at a facet F with functional λ_F such that $\lambda_F(b - a) \leq 0$. Now, c is equidistant to a and b , and a and b cannot be in the same facet of the ball centered at c (because they are in weak general position). Therefore, $\text{dist}(c, \epsilon a + (1 - \epsilon)b) < \text{dist}(c, a)$, by convexity of the ball. This contradicts the fact that $\lambda_{F_i}(b - a) \leq 0$.

Hence, we have a well-defined map $\phi : \text{bis}(a, b) \rightarrow a + H(b - a)$ given by central projection. The map ϕ is continuous since it is the restriction of central projection. It is also proper (that is, the inverse image of a compact set is compact) by a following argument: Let C be a compact subset of $H(b - a)$. By continuity, $\phi^{-1}(C)$ is closed in $\text{bis}(a, b)$, hence in $\mathbb{R}^d$ since $\text{bis}(a, b)$ itself is closed (it is the zero set of the continuous function $d(x, a) - d(x, b)$). Thus, we only need to prove that $\phi^{-1}(C)$ is bounded. This follows from the fact that

$$\phi^{-1}(C) \subset (a + \text{cone}(K)) \cap (b + \text{cone}(\overline{H(a - b)})) ,$$

and that $\text{cone}(C)$ and $\text{cone}(\overline{H(a - b)})$ are two closed linear cones meeting only at the origin, since $H(a - b)$ and $H(b - a)$ are open and disjoint in ∂K .

Once we know ϕ is proper and continuous, we only need to check that this map is bijective in order for it to be a homeomorphism. To show this, we construct its inverse. For each $v \in H(b - a)$ we consider the ray $r_v = \{a + \alpha v : \alpha \geq 0\}$. Along r_v , the distance to a is linear in α , the distance to b is convex in α and both functions are continuous. Observe also that

$$\begin{aligned} \text{dist}(a + 0v, a) &= 0, \quad \text{dist}(a + 0v, b) > 0, \\ \text{and } \lim_{\alpha \rightarrow \infty} \text{dist}(a + \alpha v, b) &< \lim_{\alpha \rightarrow \infty} \text{dist}(a + \alpha v, a) . \end{aligned}$$

The last inequality comes from the fact that as we move farther away from a along r_v , eventually $(a + \alpha v) - b$ will be in the same cone of $\mathcal{F}(K)$ as $(a + \alpha v) - a = \alpha v$ (by weak general position), and $\langle b - a, v \rangle > 0$ since $v \in H(b - a)$.

Hence, the function $\alpha \mapsto \text{dist}(a + \alpha v, b) - \text{dist}(a + \alpha v, a)$ is negative at zero, positive at infinity, continuous, and convex. Therefore, it has exactly one root, which means r_v intersects the bisector exactly once. We define $\psi(a + v)$ as this unique intersection point.

The maps ϕ and ψ are clearly inverses of one another. $\square$

Looking at the proof, the reader can check that central projection gives a proper and continuous map from $\text{bis}(a, b)$ to $a + H(b - a)$ even without assuming weak general position. We only need weak general position to construct the inverse.

Corollary 5 *If S is in weak general position and there is an empty sector $H_S(a) \in \mathcal{H}_S$ then the bisector of S is empty.*

Proof Assume that there is a point $c \in \text{bis}(S)$. For a site $a \in S$ let us show that $H_S(a) \neq \emptyset$. By definition, $c \in \text{bis}(a, b)$ for $b \in S \setminus \{a\}$. By Proposition 2, each

$\text{bis}(a, b)$ can be mapped to $H(b - a)$ by central projection. Since c is in all of these bisectors, the central projection of c into the ball $a + K$ lies in $H(b - a)$ for all b , and hence in $H_S(a)$. $\square$

The converse of Proposition 5 is true for *three* points in arbitrary dimension (Theorem 3) but not for more, even in general position, as the following example shows:

Example 2 (Empty bisector, with non-empty sectors) Consider $a = (1, -1, 0, 0)$, $b = (-1, 1, 0, 0)$, $c = (0, 0, 2, -2)$ and $d = (0, 0, -2, 2)$. Then we have

$$\begin{aligned} \text{bis}(a, b) = & \{x \mid x_3 + 1 \leq x_1, x_2, \leq x_4 - 1\} \\ & \cup \{x \mid x_4 + 1 \leq x_1, x_2, \leq x_3 - 1\} \cup \{x \mid x_1 = x_2\} . \end{aligned}$$

By symmetry, we also have

$$\begin{aligned} \text{bis}(c, d) = & \{x \mid x_1 + 2 \leq x_3, x_4, \leq x_2 - 2\} \\ & \cup \{x \mid x_2 + 2 \leq x_3, x_4, \leq x_1 - 2\} \cup \{x \mid x_3 = x_4\} . \end{aligned}$$

Since $\text{bis}(a, b, c, d)$ lies in the intersection of the two, we have

$$\text{bis}(a, b, c, d) \subseteq \{x \mid x_1 = x_2, x_3 = x_4\} .$$

So for $x \in \text{bis}(a, b, c, d)$, we may assume $x_3 = 0$, which entails:

$$\text{dist}(a, x) = \max\{x_1 + 1, 0\} - \min\{x_1 - 1, 0\} = \max\{|x_1| + 1, 2\} \leq |x_1| + 2 ,$$

with equality only when $x_1 = 0$ and

$$\text{dist}(c, x) = \max\{x_1, 2\} - \min\{x_1, -2\} = \max\{|x_1| + 2, 4\} \geq |x_1| + 2 ,$$

with equality only if $|x_1| \geq 2$. This shows that $\text{bis}(a, b, c, d) = \emptyset$.

However, $\mathcal{H}_S$ has no empty sector since the sectors of a, b, c, d contain the facets with outer normals $(1, -1, 0, 0)$, $(-1, 1, 0, 0)$, $(0, 0, 1, -1)$ and $(0, 0, -1, 1)$, respectively. Both this property and the emptiness of $\text{bis}(a, b, c, d)$ are preserved under perturbation, so this behavior happens also in general position.

3.3 Bisectors of Three Points

The goal of this section is to prove our first main result.

Theorem 3 Let $S = \{a_1, a_2, a_3\}$ be a set of three distinct points in $\mathbb{R}^d$ which lie in weak general position with respect to a polytope K . If $H_S(a_i) \neq \emptyset$ for $i = 1, 2, 3$ then $\text{bis}(a_1, a_2, a_3)$ is homeomorphic to a non-empty open subset of $\mathbb{R}^{d-2}$.

Corollary 6 For any three points in weak general position $\text{bis}(a_1, a_2, a_3)$ is either empty or pure of dimension $d - 2$. If $d = 2$ then $\text{bis}(a_1, a_2, a_3)$ is either empty or a single point.

We begin with the case $d = 2$ of Theorem 3; this occurs in [13]:

Lemma 2 *Let $S = \{a_1, a_2, a_3\} \subset \mathbb{R}^2$ be in weak general position with respect to a polytope K . If $H_S(a_i) \neq \emptyset$ for the three of them, then the bisector $\text{bis}(S)$ is a single point.*

Proof We first show that $\text{bis}(S)$ cannot contain two distinct points. To seek a contradiction, suppose that there are two points $x \neq y$ in the bisector. This means that there exist $\alpha, \beta > 0$ such that $B_x := x - \alpha K$ and $B_y := y - \beta K$ satisfy $S \subset \partial B_x \cap \partial B_y$. For brevity we call B_x and B_y *negative balls* centered at x and y , respectively. The two balls are related to one another by a homothety $\rho : B_x \rightarrow B_y$. Let $a'_i := \rho(a_i)$ and $S' = \{a'_1, a'_2, a'_3\} = \rho(S)$. Then, $S \cup S' \subset \partial B_y$, but this is impossible: the vertex sets of two homothetic triangles can only lie in the boundary of a polytope in the plane if three of the points are collinear. Three collinear points of ∂B_y necessarily lie on a single edge of B_y , and (at least) two of them would come from the same triangle S or S' , violating the weak general position of S .

It remains to exclude the case where $\text{bis}(S) = \emptyset$. Suppose that this would hold. That is, the two-point bisectors are pairwise disjoint. Now the two-point bisectors are homeomorphic to lines by Theorem 2. So the fact that they do not meet implies that each of them appears either in full or not at all in the Voronoi diagram of S . But the three of them cannot appear, since then the diagram would have four regions, not three. Thus, one of them, say $\text{bis}(a_1, a_3)$, does not appear at all in Vor_S . Consequently, $\text{bis}(a_1, a_3)$ is contained in the Voronoi region of the third point a_2 . We will show that $H_S(a_2) = \emptyset$, and this yields the desired contradiction.

To simplify the exposition, we call the line $a_1 a_3$ “horizontal”. Let u and v be the points where the ball K has a horizontal tangent, i.e., parallel to $a_1 a_3$. Without loss of generality, u is at the top of K , and v is at the bottom; see Fig. 3. The points u and v are unique and thus vertices of K , by weak general position. Consider the negative tangent cones $u + \text{cone}(u - K)$ and $v + \text{cone}(v - K)$ of K . Call K_u and K_v their respective translations that have a_1 and a_3 on the boundary, and let u' and v' be the apices of these translated cones. Such translations exist and are unique since the two boundary rays of $\text{cone}(u - K)$ point upward, and the boundary rays of $\text{cone}(v - K)$ point downward.

By construction the boundaries of K_u and K_v intersect precisely in a_1 and a_3 . Now, every negative ball of the form $B_{u',r} := u' - r(K - u)$ has K_u as its tangent cone at u' , and balls like $B_{v',r} := v' - r(K - v)$ are similar. Taking r sufficiently large, we can force

$$K_u \cap K_v = B_{u',r} \cap B_{v',r}.$$

Since $B_{u',r}$ and $B_{v',r}$ are negative balls with a_1 and a_3 in their boundaries, their centers lie on the bisector $\text{bis}(a_1, a_3)$. Our initial assumption that the bisector is contained in the Voronoi region of a_2 implies that a_2 is in the interior of both balls, i.e., in the interior of both K_u and K_v . This implies that $H_S(a_2)$ is empty: every facet-defining functional of K takes its minimum on $K_u \cap K_v$ at either a_1 or a_3 .

We conclude that, if $H_S(a_i) \neq \emptyset$ for $i = 1, 2, 3$ then $\text{bis}(S) \neq \emptyset$, and this finishes the proof. $\square$

Fig. 3 Illustrating the Proof of Lemma 2. Left: polytope K with top and bottom vertices, u and v , horizontal tangents and facet normals. Right: intersection $K_u \cap K_v$ with a_1 and a_3 on the boundary. If $a_2 \in K_u \cap K_v$, then all linear functions defined by exterior facet normals of K or, equivalently, interior facet normals of K_u and K_v , attain smaller values at a_1 or a_3 than at a_2

For the rest of the Proof of Theorem 3 let $S = \{a_1, a_2, a_3\} \subset \mathbb{R}^d$ be three points in general position with respect to a convex polytope K . The idea is to reduce the general problem to two dimensions via the following construction. Let $\text{pr} : \mathbb{R}^d \rightarrow \mathbb{R}^{d-2}$ be the affine projection that quotients out the 2-plane Π containing S . Next we exhibit relevant properties of that map.

Lemma 3 *With the above notation, let $x \in \text{int}(\text{pr}(K)) \subset \mathbb{R}^{d-2}$. Further let $\Pi_x := \text{pr}^{-1}(x)$, which is a 2-plane parallel to Π , and $K_x := K \cap \Pi_x$. Then K_x is a convex polygon.*

Proof This is a general property of polytope projections: if Q is the image of a polytope P under an affine map then the fiber of every point $x \in \text{int}(Q)$ is a polytope of dimension $\dim(P) - \dim(Q)$. $\square$

Lemma 4 *With the above notation, again let $x \in \text{int}(\text{pr}(K))$. Further, let $H_S^{(x)}(a_i)$ denote the sector of $a_i \in \Pi$ with respect to the polygon K_x . Then*

$$H_S^{(x)}(a_i) = H_S(a_i) \cap K_x.$$

Proof Note that no facet in K_x is parallel to any $a_i - a_j$ because if it were, the facet of K containing it would be parallel, too.

Let F' be a facet of K_x , and let F be the corresponding facet in K . Then, $n(F') \in \Pi_x$, the normal vector to F' , is the projection into Π of $n(F)$, and,

$$\begin{aligned}
 F' \in H_S^{(x)}(a_i) &\iff \langle n(F'), a_i \rangle > \langle n(F'), a_j \rangle \text{ for } j = 1, 2, 3 \\
 &\iff \langle n(F), a_i \rangle > \langle n(F), a_j \rangle \text{ for } j = 1, 2, 3 \\
 &\iff F \in H_S(a_i) .
 \end{aligned}$$

□

Lemma 5 *Let $S = \{a_1, a_2, a_3\}$ as before. If $H_S(a_i) \neq \emptyset$ for all i , then the intersection $\bigcap_{a_i \in S} \text{pr}(H_S(a_i)) \subset \mathbb{R}^{d-2}$ is open and not empty.*

Proof First, observe that an $x \in \partial K$ with $\text{pr}(x) \in \partial(\text{pr}(K))$ cannot be in any of the $H_S(a_i)$: indeed, $x \in \partial K$ implies that there is a normal vector of K at x orthogonal to Π , hence orthogonal to $a_i - a_j$ for every a_i, a_j . As a consequence,

$$\bigcap_{a_i \in S} \text{pr}(H_S(a_i)) \subset \text{int}(\text{pr}(K)) ,$$

which implies it is open.

For any point $x \in \text{int}(\text{pr}(K))$, the preimage $\text{pr}^{-1}(x)$ is a polygon, a slice of K . This slice has to intersect at least two of the classes of $\mathcal{H}$, because $\mathcal{H}$ is a partition (so at the slice intersects at least one class), but no class can contain a set of facets whose vectors are positively dependent (because each class is an intersection of half-spheres). Then, any point $x \in \text{int}(K)$ lies in at least two sets $\text{pr}(H_S(a_i))$.

Thus, the three open sets $\text{pr}(H_S(a_i))$ cover each point of $\text{int}(\text{pr}(K))$ at least twice. At least two of these sets must intersect, say $\text{pr}(H_S(a_1))$ and $\text{pr}(H_S(a_2))$. Suppose that $(\text{pr}(H_S(a_1)) \cap \text{pr}(H_S(a_2))) \cap \text{pr}(H_S(a_3)) = \emptyset$. Then, $\text{int}(\text{pr}(K))$ would be disconnected, because it is covered by two disjoint open sets. Since this is not possible, there must be a point in the common intersection of the three $H_S(a_i)$. □

Our next lemma finishes the Proof of Theorem 3, and it actually gives more information.

Lemma 6 *The set $\text{bis}(a_1, a_2, a_3)$ is homeomorphic to $\bigcap_{i=1,2,3} \text{pr}(H_S(a_i))$.*

Proof Consider the map

$$\phi : \text{bis}(a_1, a_2, a_3) \longrightarrow \bigcap_{i=1,2,3} \text{pr}(H_S(a_i))$$

defined as follows: Let $p \in \text{bis}(S)$, and let $v_i \in H_S(a_i)$ be the central projection from p to $a_i + \partial K$, for each $i = 1, 2, 3$. Note that each v_i lies in the corresponding $H_S(a_i)$, by Theorem 2. Further, the three points v_1, v_2 and v_3 lie in a plane parallel to Π . In particular, $\text{pr}(v_1) = \text{pr}(v_2) = \text{pr}(v_3)$ lies in $\bigcap_{i=1,2,3} \text{pr}(H_S(a_i))$ and we define

$$\phi(p) := \text{pr}(v_i) .$$

To show that ϕ is a homeomorphism, let us construct its inverse ψ . Let $\gamma : \text{pr}(\text{int}(K)) \rightarrow \text{int}(K)$ be a continuous section of pr in K . For example, but not

Fig. 4 Construction of ψ in the Proof of Theorem 3

necessarily, for each 2-plane Π' parallel to Π and intersecting K let $\gamma(\text{pr}(\Pi'))$ be the centroid of $\Pi' \cap K$.

Now, let $x \in \bigcap_{i=1,2,3} \text{pr}(H_S(a_i))$. Let $\Pi_x = \text{pr}^{-1}(x)$ and let $w_i = \gamma(x) + a_i$, for each $i = 1, 2, 3$. In the 2-plane Π_x we have a set $S_x = \{w_1, w_2, w_3\}$ and a unit ball $K \cap \Pi_x$. Lemma 3 gives that $H_{S_x}(w_i) = H_S(a_i) \cap \Pi_x$. By choice of x we have $\bigcap_i H_{S_x}(w_i) \neq \emptyset$ and Lemma 2 guarantees that the bisector of S_x is a unique point $r \in \Pi_x$; see Fig. 4.

Let v_i be the central projection of r to $w_i + \partial K_x$. Observe that $|w_i r|/|w_i v_i|$ is independent of i since it equals $\text{dist}_{K_x}(w_i, r)$, where dist_{K_x} denotes the distance induced by K_x in $\mathbb{R}^2$, and r is in $\text{bis}_{K_x}(w_1, w_2, w_3)$. Since $w_i - a_i = \gamma(x)$ is also independent of i , the three rays $a_i v_i$ meet at the point

$$p = r + \text{dist}_{K_x}(w_i, r)\gamma(x) ,$$

and $\text{dist}_K(a_i, p) = \text{dist}_{K_x}(w_i, r)$. Thus, $p \in \text{bis}(a_1, a_2, a_3)$ and we define p to be $\psi(x)$.

This gives us a well-defined map

$$\psi : \bigcap_{i=1,2,3} \text{pr}(H_S(a_i)) \longrightarrow \text{bis}(a_1, a_2, a_3) ,$$

and by construction ψ is the inverse of ϕ both ways. The map γ is continuous, as the bisector of three points in the plane depends continuously on a continuous deformation of the unit ball. So ψ is continuous. Thus, ϕ and ψ are homeomorphisms between $\text{bis}(a_1, a_2, a_3)$ and $\bigcap_{i=1,2,3} \text{pr}(H_S(a_i))$. Since the latter is not empty and open by Lemma 5, the former is homeomorphic to a non-empty open subset of $\mathbb{R}^{d-2}$. $\square$

Our Proof of Theorem 3 closely follows the proof of the 3-dimensional case in [14]. There, it is additionally shown that the number of connected components of $\text{bis}(a_1, a_2, a_3)$ equals the total number of connected components of the three sectors $H_S(a_i)$ minus two. We can extend this formula to higher dimension and to higher Betti numbers (over any field):

Theorem 4 Let $a_1, a_2, a_3 \in \mathbb{R}^d$ be three points in weak general position with respect to a polytope K and assume that $H_S(a_i) \neq \emptyset$ for all three. Then, for $j \in \{0, \dots, d-3\}$, we have

$$\tilde{\beta}_j(\text{bis}_K(a_1, a_2, a_3)) = \sum_{i=1}^3 \tilde{\beta}_j(H_S(a_i)) .$$

Proof Consider the same projection $\text{pr} : \mathbb{R}^d \rightarrow \mathbb{R}^{d-1}$ as before. Observe that, since $\Pi_x \cap H_S(a_i)$ is empty or contractible for every plane Π_x parallel to pr , we have

$$H_S(a_i) \simeq \pi(H_S(a_i)) \quad \text{for } i = 1, 2, 3 .$$

We now apply Alexander duality in the one-point compactification $\mathcal{S}$ of $\text{int}(\text{pr}(K))$, which is a sphere of dimension $d-2$. Alexander duality says that if U is an open and locally contractible subset of $\mathcal{S}$ then

$$\tilde{\beta}_j(U) = \tilde{\beta}_{d-3-j}(\mathcal{S} \setminus U) .$$

In particular, if we let $C_i = \mathcal{S} \setminus \text{pr}(H_S(a_i))$ we have

$$\sum_{i=1}^3 \tilde{\beta}_j(\text{pr}(H_S(a_i))) = \sum_{i=1}^3 \tilde{\beta}_{d-3-j}(C_i) ,$$

and

$$\tilde{\beta}_j\left(\bigcap_{i=1}^3 \text{pr}(H_S(a_i))\right) = \tilde{\beta}_{d-3-j}\left(\bigcup_{i=1}^3 C_i\right) .$$

Yet C_1, C_2 and C_3 are pairwise disjoint except for the “point at infinity” of $\mathcal{S}$, because each point of $\text{int}(\text{pr}(K))$ lies in at least two of the sets $\text{pr}(H_S(a_i))$. Thus, $\bigcup_{i=1}^3 C_i$ is the topological wedge (or one-point sum) of C_1, C_2 and C_3 , which makes the right-hand sides of the two last equations coincide. $\square$

Remark 6 One may ask how complicated the Betti numbers $\tilde{\beta}_j(H_S(a_i))$ in Theorem 4 can be. Equivalently, how complicated the topology of three point bisectors can be. Such a bisector is $(d-2)$ -dimensional, so the relevant Betti numbers are $\tilde{\beta}_0, \dots, \tilde{\beta}_{d-2}$. The last one, $\tilde{\beta}_{d-2}$, must vanish as $\tilde{\beta}_j(H_S(a_i)) = \tilde{\beta}_j(\text{pr}(H_S(a_i)))$, and the latter is an open subset of $\mathbb{R}^{d-2}$. But $\tilde{\beta}_{d-3}$ can be nonzero, as the following example shows.

Example 3 Let $K = \mathbb{B}^3$ be the tropical unit ball, i.e., $d = 3$. The three sites $a = (0, 0, 4)$, $b = (-3, 0, 2, 0)$ and $c = (0, -3, 0, 2)$ lie in weak general position in $\mathbb{R}^4/\mathbb{R}\mathbf{1}$. We describe each sector in terms of the facets whose relative interiors lie in that sector. The facet F_{ij} is the one at which coordinate i is minimized and j is maximized. We obtain

$$\begin{aligned} H(a) &= (F_{14}, F_{23}) , & H(b) &= (F_{12}, F_{13}, F_{32}, F_{42}, F_{43}) , \\ H(c) &= (F_{21}, F_{24}, F_{31}, F_{34}, F_{41}) . \end{aligned}$$

Fig. 5 Tropical bisectors for $b - a = (-1, 1, 0)$, $(-1, 1, \frac{1}{2})$ and $(-1, 1, 1)$, respectively. The first two are in weak general position but the last one is not. Only the middle one is in general position. The bisectors are shown in black and the rest of the tropical hypersurface containing it in gray; see Proposition 4

The sector $H(a)$ is disconnected; i.e., $\tilde{\beta}_0(H(a))$ is nonzero. By Theorem 4 $\text{bis}(a, b, c)$ is disconnected, too.

Remark 7 The results in this section and in Sect. 3.2 involve only *weak* general position. While we keep assuming that K is a polytope for simplicity, these results generalize to an arbitrary convex body K . In that case S is in weak general position if the boundary ∂K contains no segment parallel to a difference $a - b$ with $a, b \in S$.

4 Classification of Tropical Bisectors of Two Points

4.1 Tropical Bisectors and Tropical Hypersurfaces

In the classical case the bisector of two points is a degenerate quadric, namely the affine hyperplane perpendicular to the connecting line segment and which runs through the midpoint. The tropical analog is more interesting.

Proposition 4 *Let $a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$ be in weak general position. Then the homogeneous max-tropical Laurent polynomial*

$$\phi(a, b) = \max \left(\max_{i, j \in [d+1]} (x_i - a_i - x_j + a_j), \max_{k, \ell \in [d+1]} (x_k - b_k - x_\ell + b_\ell) \right) \quad (6)$$

vanishes on the bisector $\text{bis}(a, b)$. That is, the set $\text{bis}(a, b)$ is contained in a max-tropical hypersurface of degree $d + 1$.

Proof Recall that a max-tropical (Laurent) polynomial vanishes if the maximum is attained at least twice; see [22, §3.1]. First, we check that there are no duplicates among the terms in the representation (6) of $\phi(a, b)$. Assume the contrary, i.e., $x_i - a_i - x_j + a_j = x_i - b_i - x_j + b_j$ for some $i, j \in [d + 1]$. Then $a_j - a_i = b_j - b_i$, which forces $\langle e_j - e_i, b - a \rangle = (b_j - a_j) - (b_i - a_i) = 0$. Thus $b - a$ is parallel

to the facet of $\mathbb{B}^d$ with normal vector $e_j - e_i$. This was explicitly excluded in our assumption, and we arrive at the desired contradiction. We infer that the $2d(d+1)$ terms are pairwise distinct.

Let $x \in \text{bis}(a, b)$. This means that $\text{dist}(a, x) = \text{dist}(b, x)$, and thus

$$\max_{i, j \in [d+1]} (x_i - a_i - x_j + a_j) = \max_{k, \ell \in [d+1]} (x_k - b_k - x_\ell + b_\ell).$$

It follows that $\phi(a, b)$ vanishes at x .

The degree of the bisector tropical hypersurface can be read off any Laurent monomial like $x_i - x_j$ by adding $x_1 + x_2 + \dots + x_{d+1}$, which yields the true monomial $2x_i + \sum_{k \in [d+1] - \{i, j\}} x_k$. The latter has degree $2 + (d+1-2) = d+1$. $\square$

Proposition 4 yields a trivial algorithm to compute tropical bisectors in weak general position: enumerate the maximal cells of the tropical hypersurface defined by (6) and select those maximal cells that attain maxima in one monomial of type $x_i - a_i - x_j + a_j$ and one monomial of type $x_k - b_k - x_\ell + b_\ell$. This algorithm needs to go through the $\Theta(d^4)$ choices of one monomial from the left and one from the right. This is worst case optimal, as we will prove in Corollary 8 that tropical bisectors can have $\Omega(d^4)$ maximal cells.

Example 4 The labeling of the faces of a tropical bisector does not need to be unique if a and b are not in weak general position. For instance, if $b - a = (-1, 1, 0)$ then

$$\begin{aligned} \text{bis}_{(-+*), (+-*)}(a, b) &= \text{bis}_{(-++), (+--)}(a, b) \\ &= \text{bis}_{(-++), (+--)}(a, b) = \text{bis}_{(-+-), (+-+)}(a, b) \\ &= \text{bis}_{(-+-), (+-+)}(a, b) \end{aligned}$$

is the only face; see Fig. 5 (left).

4.2 The Bisection Fan

Normal equivalence of tropical bisectors is preserved by translation and scaling. In particular the equivalence class of $\text{bis}(a, b)$ is uniquely determined by the direction of the vector $b - a$. The *bisection fan* $\mathcal{F}_{\text{bis}}^d$ is the complete polyhedral fan in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ whose relatively open cones are defined by “ x and y lie in the same cone if and only if $\text{bis}(0, x)$ and $\text{bis}(0, y)$ are normally equivalent”. Put differently, two points $a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ are in *general position* if and only if the difference $b - a$ lies in a maximal cone of $\mathcal{F}_{\text{bis}}^d$. In the rest of this section we show that $\mathcal{F}_{\text{bis}}^d$ is indeed a polyhedral fan and give an explicit description of it.

Recall that an *ordered partition* or *total preorder* on a finite set S is a partition of S into non-empty parts together with a total order on the parts. If the parts are denoted $S_1, \dots, S_k$ (in this order), we can write $x \leq y$ meaning “ $x \in S_i$ and $y \in S_j$ for some $i \leq j$ ”. In particular, for all $x, y \in S$ we have $x \leq y \leq x$ if and only if x and y lie in the same part.

Any real vector $v = (v_1, \dots, v_{d+1}) \in \mathbb{R}^{d+1}$ induces an ordered partition $S(v)$ of $[d+1]$ by putting together the coordinates that have the same value and ordering the

Fig. 6 The fans $\mathcal{F}(\mathbb{B}^3)$, $\mathcal{F}(A_3)$, and the bisection fan $\mathcal{F}_{\text{bop}}^3$

groups according to their values. For example, the vector $v = (3, 1, 6, 4, 6, 3, 1)$ of length seven induces the partition $(\{2, 7\}, \{1, 6\}, \{4\}, \{3, 5\})$ of the set $\{1, 2, \dots, 7\}$ into four parts. Note that the ordered partition $S(v)$ is constant on the class $v + \mathbb{R}\mathbf{1}$. Hence these ordered partitions are defined for points in the projective tropical torus $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$.

For $v, w \in \mathbb{R}^{d+1}$ with $S(v) = S(w)$ we have $S(v + w) = S(v)$ and, moreover, $S(\alpha v) = S(v)$ for any positive real α . That is to say, the stratification of $\mathbb{R}^{d+1}$ by ordered partitions forms a complete polyhedral fan. In what follows we seek to refine that fan by recording which part, or which gap between parts, contains the *midvalue*

$$\mu(v) := \frac{1}{2} \left(\max_{i \in [d+1]} v_i + \min_{i \in [d+1]} v_i \right).$$

The *bisected ordered partition* of $[d + 1]$ induced by $v \in \mathbb{R}^{d+1}$ is the ordered partition $S(v)$ as defined above, together with the information of which (values of) parts are smaller, equal or greater than the midvalue $\mu(v)$; see also [9]. Equivalently, this is the ordered partition associated with the *extended vector* $(v, \mu(v)) \in \mathbb{R}^{d+2}$. We denote by $\mathcal{F}_{\text{bop}}^d$ the fan of bisected ordered partitions of dimension d .

Remark 8 The “finest” or “most generic” ordered partitions are the permutations, in which each part is a singleton. Hence, the fan of ordered partitions equals the normal fan of the *permutahedron*. This, in turn, coincides with the fan of regions in the *braid arrangement* or Coxeter arrangement of type A_d , which consists of the hyperplanes $\{x \mid x_i = x_j\}$ for $1 \leq i < j \leq d + 1$. We denote this fan $\mathcal{F}(A_d)$. It is intermediate between the central fan of the tropical ball (which is coarser) and the fan of bisected ordered partitions (which is finer):

$$\mathcal{F}(\mathbb{B}^d) \leq \mathcal{F}(A_d) \leq \mathcal{F}_{\text{bop}}^d;$$

see Fig. 6 for a visualization of the case $d = 3$. Note that $\mathcal{F}(A_d)$ is also the fan of weak general position: a and b are in weak general position if and only if $b - a$ lies in a full-dimensional cell.

Fig. 7 The bisection fan, $\mathcal{F}_{\text{bis}}^d$, for $d = 2$. The three vectors $b - a$ for Fig. 5 have been marked

Example 5 In $d = 2$, the fans $\mathcal{F}(A_d)$ and $\mathcal{F}(\mathbb{B}^d)$ coincide, they form the face fan of the regular hexagon $\mathbb{B}^2$. The bisection fan $\mathcal{F}_{\text{bop}}^d$ is the barycentric subdivision of it; see Fig. 7. Excluding permutations of the coordinates, and sign inversion, we infer that there are three types of tropical bisectors in the plane, and these are shown in Fig. 5. The type to the left is in weak general position but not in general position, the type to the left is in general position, and the type to the right is not even in weak general position.

Recall that the max-tropical *line segment* spanned by two points, a and b , is the set

$$[a, b] := \{\max(\alpha \mathbb{1} + a, \beta \mathbb{1} + b) \mid \alpha, \beta \in \mathbb{R}\} \subset \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}.$$

It is worth noting that the combinatorial types of tropical segments are classified by the braid fan $\mathcal{F}(A_d)$; see [22, Prop. 5.11]. In this sense the fan of bisected ordered partitions $\mathcal{F}_{\text{bop}}^d$ classifies the combinatorial types of “bisected tropical segments”:

Proposition 5 *The bisected ordered partition of $b - a$ contains the same information as the combinatorial type of the tropical line segment $[a, b]$ together with the information of which part contains the midpoint.*

Proof Suppose for simplicity that $a, b \in \mathbb{R}^{d+1}$ satisfy

$$b_1 - a_1 \leq b_2 - a_2 \leq \dots \leq b_{d+1} - a_{d+1}. \quad (7)$$

Then $[a, b]$ is the union of at most d ordinary line segments, one for each subset of coordinates between a strict inequality in (7). That is, the combinatorics of the tropical

segment is the same as the ordered partition of $b - a$. The midvalue $\mu(b - a)$ selects one of the ordinary segments. $\square$

The goal of the rest of this section is to prove the following:

Theorem 5 *We have $\mathcal{F}_{bis}^d = \mathcal{F}_{bop}^d$. That is, given $a, b, a', b' \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ the polyhedra $\text{bis}(a, b)$ and $\text{bis}(a', b')$ are normally equivalent if and only if $b - a$ and $b' - a'$ induce the same bisected ordered partition of $[d + 1]$.*

4.3 Proof of Theorem 5

In (3) we defined the type (F_-, F_*, F_+) of a face F of $\mathbb{B}^d$ or the corresponding face cone in $\mathcal{F}(\mathbb{B}^d)$. For a pair of faces, F and G , this gives rise to the following *labeling partition* of $[d + 1]$:

$$\begin{aligned} L_0 &:= (F_- \cap G_-) \cup (F_+ \cap G_+) , \\ L_+ &:= (F_+ \cap G_*) \cup (F_* \cap G_-) , \\ L_- &:= (F_- \cap G_*) \cup (F_* \cap G_+) , \\ L_{+1} &:= F_+ \cap G_- , \\ L_{-1} &:= F_- \cap G_+ , \\ L_* &:= F_* \cap G_* . \end{aligned} \tag{8}$$

As a first step in the Proof of Theorem 5, the following lemma characterizes when $\text{bis}_{(F,G)}(a, b)$ is non-empty. Recall that this is the case if and only if there is a tropical ball touching a and b at faces F and G , respectively.

Lemma 7 *Let F and G be a fixed pair of faces of $\mathbb{B}^d$ with the labeling partition defined as in (8). Further let $a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. Then the set $\text{bis}_{(F,G)}(a, b)$ is not empty if and only if there exist $\gamma \in \mathbb{R}$ and $\delta \in [0, \infty)$ such that the following conditions are satisfied:*

$$\left\{ \begin{array}{ll} (b - a)_i = \gamma & \text{if } i \in L_0 , \\ (b - a)_i \in [\gamma, \gamma + \delta] & \text{if } i \in L_+ , \\ (b - a)_i \in [\gamma - \delta, \gamma] & \text{if } i \in L_- , \\ (b - a)_i = \gamma - \delta & \text{if } i \in L_{-1} , \\ (b - a)_i = \gamma + \delta & \text{if } i \in L_{+1} , \\ (b - a)_i \in [\gamma - \delta, \gamma + \delta] & \text{if } i \in L_* . \end{array} \right. \tag{9}$$

Proof Let us assume that the face of the tropical bisector $\text{bis}(a, b)$ defined by (F, G) is non-empty. Then there is a point x such that $\text{dist}(a, x) = \text{dist}(x, b) = \delta$ and $a - x \in F$ as well as $b - x \in G$. We set $\gamma_a = \min_{i \in [d+1]} (a_i - x_i)$ and $\gamma_b = \min_{i \in [d+1]} (b_i - x_i)$. The possible values for the coordinates of $a - x$ and $b - x$ are

$$\left\{ \begin{array}{ll} (a - x)_i = \gamma_a + \delta & \text{if } i \in F_- , \\ (a - x)_i = \gamma_a & \text{if } i \in F_+ , \\ (a - x)_i \in [\gamma_a, \gamma_a + \delta] & \text{if } i \in [d + 1] \setminus (F_- \cup F_+) \end{array} \right. \tag{10}$$

and

$$\begin{cases} (b-x)_i = \gamma_b + \delta & \text{if } i \in G_- , \\ (b-x)_i = \gamma_b & \text{if } i \in G_+ , \\ (b-x)_i \in [\gamma_b, \gamma_b + \delta] & \text{if } i \in [d+1] \setminus (G_- \cup G_+) , \end{cases} \quad (11)$$

for some $\delta \geq 0$. Setting $\gamma = \gamma_b - \gamma_a$ the above translates into (9).

For the converse, note that going from (10) and (11) to (9) is the Fourier-Motzkin elimination of the variables x_i . Therefore, any γ and $\delta \geq 0$ which are feasible for (9) can be lifted to a solution of (10) and (11). That is to say, we can set $\gamma_a = 0$ and $\gamma_b = \gamma$, and the conditions in (10) and (11) yield a point $x \in \text{bis}_{(F,G)}(a, b)$. $\square$

Proposition 6 *Let F and G be a fixed pair of faces of $\mathbb{B}^d$. Then the set*

$$C := \left\{ b - a \mid a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1} \text{ with } \text{bis}_{(F,G)}(a, b) \neq \emptyset \right\} \quad (12)$$

is both a polyhedral cone and a tropical cone, although perhaps not a tropical polyhedral cone.

Proof Let $a, a', b, b' \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ with $\text{bis}_{(F,G)}(a, b)$ and $\text{bis}_{(F,G)}(a', b')$ both non-empty. Since $\text{bis}_{(F,G)}(a, b) \neq \emptyset$, by Lemma 7, there are scalars γ and δ satisfying the conditions (9). Likewise there are certificates γ' and δ' for $\text{bis}_{(F,G)}(a', b') \neq \emptyset$. By linearity of the conditions (9) it follows that $\gamma + \gamma'$ and $\delta + \delta'$ certify that $\text{bis}_{(F,G)}(a + a', b + b') \neq \emptyset$: for instance, we have $(b + b' - a - a')_i = \gamma + \gamma'$ for $i \in L_0 = (F_- \cap G_-) \cup (F_+ \cap G_+)$. Since clearly $\alpha c \in C$ for all $c \in C$ and $\alpha \geq 0$ we conclude that C is an ordinary cone. This cone is polyhedral as it is defined in terms of the finitely many linear conditions (9).

A similar argument shows that C is also closed with respect to taking arbitrary $(\max, +)$ -linear combinations: for instance, with the above notation we have $\max((b - a)_i, (b' - a')_i) = \max(\gamma, \gamma')$ for $i \in L_0$. This shows that C is a tropical cone. $\square$

Corollary 7 *The bisection fan of tropical bisectors is a classical polyhedral fan, and a tropical (perhaps not tropical polyhedral) fan.*

Proof We know that the feasibility region of a face (F, G) is a tropical and classical polyhedral cone. Finite intersections of these cones are again tropical cones, classical cones, and polyhedral cones. Therefore, the feasibility region of a normal equivalence class, which is the intersection of the cones of its non-empty faces, is again a tropical and classical polyhedral cone. Hence, the whole fan has this structure. $\square$

The following shows one direction of Theorem 5, namely, $\mathcal{F}_{\text{bis}}^d$ is coarser than $\mathcal{F}_{\text{bop}}^d$.

Lemma 8 *Let $F, G \in \mathcal{F}(\mathbb{B}^d)$. Then, whether $\text{bis}_{(F,G)}(a, b)$ is empty or not, for each $a, b \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ depends only on the bisected ordered partition of $b - a$.*

Proof Consider the partition of $[d+1]$ into six sets $L_0, L_+, L_-, L_{+1}, L_{-1}, L_*$ defined in (8). We want to show that feasibility of the system (9) for a given a and b depends

only on the bisected ordered partition of $b - a$. Without loss of generality we assume $a = 0$ and $b_1 \leq b_2 \leq \dots \leq b_{d+1}$ as in (7).

Let $\mu(b) := \frac{1}{2}(b_{d+1} + b_1) = \frac{1}{2} \text{dist}(0, b) + b_1$ be the midvalue of $b - 0$ and $m = \max(0 + \mu(b)\mathbb{1}, b)$ the midpoint of the segment $[0, b]$. In particular, $\text{dist}(0, m) = \text{dist}(m, b) = \frac{1}{2} \text{dist}(0, b)$.

We distinguish three cases, depending on whether both, none, or exactly one of L_{+1} and L_{-1} are empty.

Claim I: Suppose $L_{+1} \cup L_{-1} = \emptyset$. Then, (9) is feasible if and only if there are $k \leq \ell$ with

$$L_- \subseteq \{1, \dots, k\} \quad L_0 \subseteq \{k+1, \dots, \ell\}, \quad L_+ \subseteq \{\ell+1, \dots, d+1\}. \quad (\text{I.1})$$

Indeed, in this case feasibility of (9) is equivalent to feasibility of

$$\begin{cases} \gamma \geq b_i & \text{for } i \in L_- , \\ \gamma = b_i & \text{for } i \in L_0 , \\ \gamma \leq b_i & \text{for } i \in L_+ , \end{cases} \quad (13)$$

which implies the ordered partition to satisfy (I.1). Conversely, if the ordered partition satisfies (I.1), then let γ be chosen to satisfy (13) and let $\delta \geq \min(\gamma - b_1, b_{d+1} - \gamma)$. This yields a feasible solution to (9). In other words, in this case we can tell if C is empty or not by just looking at the the ordered partition of b ; the relative position of the midpoint is irrelevant.

Claim II: Suppose $L_{+1} \neq \emptyset = L_{-1}$. Then, (9) is feasible if and only if, in addition to (I.1), we have

$$\{b_i \mid i \in L_{+1}\} = \{b_{d+1}\}, \quad (\text{II.1})$$

$$b_i \leq \mu(b) \text{ for } i \in L_0 \cup L_- , \quad (\text{II.2})$$

$$|\{b_i \mid i \in L_0\}| \leq 1. \quad (\text{II.3})$$

If (I.1), (II.1), (II.2), and (II.3) hold, take $\gamma = \max\{b_i \mid i \in L_0 \cup L_- \cup \{1\}\}$ and $\delta = b_{d+1} - \gamma$.

Conversely, if (γ, δ) is feasible for (9) then $\gamma + \delta = b_i = b_{d+1}$ for all $i \in L_{+1}$. In particular, $\text{bis}_{(F,G)}(a, b)$ is empty unless $\{b_i \mid i \in L_{+1}\} = \{b_{d+1}\}$ is a singleton. Since the coefficients of b are in ascending order, it follows that

$$\gamma + \delta = b_{d+1} \quad (14)$$

and $b_i = b_{d+1}$ for all $i \in L_{+1}$. This shows that (II.1) holds. Now the constraints of (9) translate into (13) as in the previous case, which implies (I.1). Additionally

$$\gamma - \delta \leq b_1. \quad (15)$$

Adding (14) and (15) now yields

$$\gamma \leq \frac{1}{2}(b_1 + b_{d+1}) = \mu(b) ,$$

which, by (13), gives (II.2). Finally, (II.3) follows from the fact that the only possible value in $\{b_i \mid i \in L_0\}$ is γ .

The case where $L_{-1} \neq \emptyset$ and $L_{+1} = \emptyset$ is analogous.

Claim III: Suppose $L_{+1} \neq \emptyset \neq L_{-1}$. Then, (9) is feasible if and only if

$$\{b_i \mid i \in L_{-1}\} = \{b_1\} \quad \{b_i \mid i \in L_{+1}\} = \{b_{d+1}\} , \quad (\text{III.1})$$

$$b_i \leq \mu(b) \text{ for } i \in L_{-} , \quad b_i = \mu(b) \text{ for } i \in L_0 , \quad b_i \geq \mu(b) \text{ for } i \in L_{+} . \quad (\text{III.2})$$

Indeed, in this case the only candidate solution for (9) is $\gamma = \mu(b) = (b_{d+1} + b_1)/2$ and $\delta = (b_{d+1} - b_1)/2$. This is a solution or not depending only on whether Eqs. (III.1) and (III.2) are satisfied. $\square$

The following gives the other direction of Theorem 5: $\mathcal{F}_{\text{bis}}^d$ refines $\mathcal{F}_{\text{bop}}^d$.

Lemma 9 Let $a, b, a', b' \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ be two pairs of points. If the bisected ordered partitions of $b - a$ and $b' - a'$ are not the same, then there is a pair of faces $F, G \in \mathcal{F}(\mathbb{B}^d)$ such that exactly one of $\text{bis}_{F,G}(a, b)$ or $\text{bis}_{F,G}(a', b')$ is empty.

Proof As before, we assume without loss of generality that $a = a' = 0$. We know that the bisected ordered partitions of b and b' are different. Our goal is to find a pair of faces (F, G) that lies in one and only one of the bisectors. We do this in three cases, depending on the difference between the bisected ordered partitions

Case I: Suppose that

$$\begin{aligned} \left\{ i \in [d+1] \mid b_i = \max_j b_j \right\} &\neq \left\{ i \in [d+1] \mid b'_i = \max_j b'_j \right\} \text{ or} \\ \left\{ i \in [d+1] \mid b_i = \min_j b_j \right\} &\neq \left\{ i \in [d+1] \mid b'_i = \min_j b'_j \right\} . \end{aligned}$$

Without loss of generality, there is an index i such that b_i is maximum and b'_i is not, or b_i is minimum and b'_i is not. Let F be the face with

$$F_{+} = \left\{ i \in [d+1] \mid b_i = \max_j b_j \right\} , \quad F_{-} = \left\{ i \in [d+1] \mid b_i = \min_j b_j \right\} ,$$

and let $G = -F$. This choice makes

$$L_{+1} = F_{+} , \quad L_{-1} = F_{-} , \quad L_{*} = F_{*} , \quad L_{-} = L_{+} = L_0 = \emptyset .$$

The cell $\text{bis}_{F,G}(a, b)$ is not empty by Lemma 7, since the following is a solution for (9):

$$\begin{aligned}\gamma &= \frac{1}{2} \left(\max_{i \in [d+1]} b_i + \min_{i \in [d+1]} b_i \right), \\ \delta &= \frac{1}{2} \left(\max_{i \in [d+1]} a_i - \min_{i \in [d+1]} b_i \right) = \frac{1}{2} \text{dist}(a, b).\end{aligned}$$

However, $\text{bis}_{(F,G)}(a', b')$ is empty: in order for it not to be empty we would need

$$\left\{ i \in [d+1] \mid b'_i = \min_j b'_j \right\} \subset F_- , \quad \left\{ i \in [d+1] \mid b'_i = \max_j b'_j \right\} \subset F_+ .$$

Case II: Suppose that b and b' have exactly the same maxima and minima but the ordered partitions of b and b' do not coincide. That is, there is a pair of indices, $i, j \in [d+1] \setminus \{1, d+1\}$ such that $b_i \geq b_j$ but $b'_i < b'_j$.

We may assume that 1 and $d+1$ are a minimum and a maximum, respectively, of both b and b' . Let F be the face with $F_+ = \{d+1\}$ and $F_- = \{1\}$. Let G be the face with $G_+ = \{i\}$, and $G_- = \{j\}$. Then, (8) gives

$$\begin{aligned}L_+ &= \{d+1, j\}, & L_- &= \{1, i\}, \\ L_* &= [d+1] \setminus (L_+ \cup L_-), & L_{-1} &= L_{+1} = L_0 = \emptyset.\end{aligned}$$

Then, $\text{bis}_{F,G}(a, b)$ is not empty since $\gamma = (b_i + b_j)/2$, and $\delta = \text{dist}(a, b)$ is a solution of (9). However, the system for b' is infeasible, since $b'_i < b'_j$. Thus, $\text{bis}_{F,G}(a', b')$ is empty.

Case III: Suppose that b and b' have exactly the same maxima and minima and the same ordered partitions but the midvalue does not coincide.

As before, we assume without loss of generality that 1 and $d+1$ are a minimum and a maximum, respectively, of both b and b' . Then, there is an index $i \in [d+1] \setminus \{1, d+1\}$ such that $\mu(b) \leq b_i$ but $\mu(b') > b'_i$ (or vice-versa, but that would give an equivalent case).

In this case, we let F and G be the faces with $F_+ = \{d+1\}$, $F_- = \{1\}$, $G_+ = \{1\}$ and $G_- = \{i\}$. These faces produce

$$\begin{aligned}L_+ &= \{d+1, i\}, & L_{-1} &= \{1\}, \\ L_* &= [d+1] \setminus (L_+ \cup L_{-1}), & L_- &= L_{+1} = L_0 = \emptyset.\end{aligned}$$

Then, $\text{bis}_{F,G}(a, b)$ is not empty since $\gamma = \mu(b)$, $\delta = \text{dist}(a, b)/2$ is a solution of (9) for b

However, the system for b' is infeasible. This is because (9) specifies that $\gamma + \delta \geq b'_{d+1}$, and $\gamma - \delta = b'_1$. Adding them together and dividing by two we get $\gamma \geq (b'_1 + b'_{d+1})/2 = \mu(b')$. We also need by (9) and $i \in L_+$ that $\gamma \leq b'_i$. Then, $\mu(b) \leq b'_i$, which contradicts our assumption. $\square$

Corollary 8 *The tropical bisector of two points in general position has $\Theta(d^4)$ maximal cells.*

Proof The upper bound is trivial, each maximal cell corresponds to a choice of a pair of facets (F, G) from the tropical ball. For the lower bound, assume without loss of generality that $a = 0$ and $b_1 < b_2 < \dots < b_{d+1}$. Then, for each choice of $i, j, k, \ell \in \{1, \dots, d+1\}$, all different and with $\max\{j, \ell\} < \min\{i, k\}$, let $F_+ = \{i\}$, $F_- = \{j\}$, $G_+ = \{k\}$, $G_- = \{\ell\}$. By Claim 1 in the Proof of Lemma 8 the set $\text{bis}_{(F,G)}(a, b)$ is not empty. Since a and b are in general position and F and G are facets of $\mathbb{B}^d$, we have $\dim \text{bis}_{(F,G)}(a, b) = d - 1$. There are $4\binom{d+1}{4}$ ways of choosing such $\{i, j, k, \ell\}$. $\square$

4.4 The Structure of Tropical Voronoi Regions

A *polytrope* is an ordinary polytope which is also convex in the tropical sense (with respect to min and max simultaneously); see [17]. These are precisely the ordinary polytopes whose facets normals are roots of type A_d , i.e., $e_i - e_j$ for $i \neq j$; they generalize the “alcoved polytopes” of Lam and Postnikov [20]. Here we relax this notion by also calling a not necessarily bounded ordinary polyhedron a *polytrope* if its facets normals are roots of type A_d ; this was called a “weighted digraph polyhedron” in [18].

The tropical unit ball $\mathbb{B}^d$ is a polytrope. But a more important example for us are the polytropes $Q = \bigcap_{a \in S} (a_+ F_a)$, where $F_a \in \mathbb{B}^d$ for each $a \in S$. Recall from Sect. 3 that in such a Q bisectors of subsets of S agree with affine subspaces. Thus:

Lemma 10 *For each polytrope Q as above and $a \in S$, the set $Q \cap \text{Vor}_S(a)$ is the intersection of Q with ordinary affine halfspaces with facet normal $e_i - e_j - e_k + e_\ell$, where i and j are fixed.*

Proof Let i and j be the coordinates maximized and minimized in F_a , respectively. For each $b \in S \setminus a$, the condition for x to be closer to a than to b is that $x_i - a_i - x_j + a_j \leq x_k - b_k - x_\ell + b_\ell$, where k and ℓ are the coordinates corresponding to F_b ; see Proposition 4. $\square$

We call the intersection of a (possibly unbounded) polytrope with ordinary affine halfspaces with facet normal $e_i - e_j - e_k + e_\ell$, where i and j are fixed, a *semi-polytrope* of type (i, j) . A semi-polytrope in $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1} \cong \mathbb{R}^d$ has at most $2\binom{d+1}{2}$ facets, since there are at most $(d+1)d$ vectors $e_i - e_j - e_k + e_\ell$ for $k \neq \ell$ and fixed (i, j) , plus the (at most) $(d+1)d$ facets of a polytrope.

A set $X \subset \mathbb{R}^{d+1}/\mathbb{R}\mathbf{1} \cong \mathbb{R}^d$ is *star convex* with center c if for any point $x \in X$ the ordinary line segment $[c, x]$ is contained in X . Clearly any convex set is star convex, but the converse does not hold. Star convex sets are contractible. Despite the many differences to Euclidean Voronoi diagrams, the following result expresses a key similarity.

Theorem 6 *Let $S \subset \mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$ be a finite set in weak general position. Then each tropical Voronoi region of S is the star convex union of finitely many (possibly unbounded) semi-polytropes.*

Fig. 8 Tropical Voronoi diagram of five points in $\mathbb{R}^3/\mathbb{R}1$. The decomposition of Voronoi regions into semi-polytopes is shown by dashed lines

Proof That Voronoi regions for polyhedral norms are star-convex is a well-known fact (see [4, p. 133] or [23, p. 127]), which follows for example from Theorem 2. By Lemma 10, $\text{Vor}_S(a)$ decomposes as finitely many semi-polytopes, by intersecting it with the individual polyhedra $Q = \bigcap_{a \in S} (a + F_a)$, for all choices of $\{F_a\}_{a \in S}$. $\square$

Semi-polytopes are not necessarily tropically convex, and this entails that the regions of a tropical Voronoi diagram are not necessarily tropically convex either; see Fig. 8 for an example. The tropical torus $\mathbb{R}^{d+1}/\mathbb{R}1$ is compactified by the tropical projective space $\mathbb{TP}^d$; the latter is the max-tropical convex hull of the $d + 1$ max-tropical unit vectors

$$(0, -\infty, -\infty, \dots, -\infty), (-\infty, 0, -\infty, \dots, -\infty), \dots, (-\infty, -\infty, \dots, -\infty, 0).$$

In this way, $\mathbb{TP}^d$ may be seen as an infinitely scaled tropical unit ball, which is a polytope; see [18, §3.5]. Similarly for arbitrary (semi-)polytopes the line between bounded and unbounded is blurred in the compactification.

5 Computing Tropical Voronoi Diagrams

We will discuss several algorithms. Some of these methods are similar to their classical Euclidean counterparts, others rely on tailored data structures, which are based on Theorem 6. For the complexity analysis of our algorithms we will consider the dimension as constant.

5.1 The Planar Case

There are several methods for computing Euclidean Voronoi diagrams in $\mathbb{R}^2$ with the optimal time complexity $O(n \log n)$ and linear space; see [7, §7.2]. This agrees with the situation for planar tropical convex hull computations; see [16, §5]. Chew and Drysdale [6] gave a divide-and-conquer algorithm with the same complexity for planar Voronoi diagrams with respect to arbitrary norms. Here we sketch a tropical analog of Fortune's beach line algorithm [10]; see also [27].

Suppose that we are given a set S of n sites in $\mathbb{R}^3/\mathbb{R}\mathbb{1}$. In view of Theorem 6 the tropical Voronoi diagram of S gives rise to a planar graph where vertices are circumcenters of triples of points in S , edges are two point bisectors, and faces are Voronoi regions. We can make this planar embedding piecewise linear by subdividing each bisector into at most five segments; see Fig. 5. The relevant data structure, as in the classical setting, is a doubly-connected edge list which requires $O(n)$ space; see [7, §2.2].

The beach line algorithm is based on a line sweep. The *tropical sweep line* at time t in $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ is the set $L(t) = (0, t, 0) + \mathbb{R}(0, 0, 1) + \mathbb{R}\mathbb{1}$. Note that $L(t)$ is an ordinary line which is also tropically convex (with respect to min and max). For an arbitrary point x we call the set

$$P(x, t) = \left\{ a \in \mathbb{R}^3/\mathbb{R}\mathbb{1} \mid \text{dist}(x, a) = \text{dist}(x, L(t)) \right\}$$

the *parabola spanned by x and $L(t)$* ; here $\text{dist}(x, L(t)) = \min\{\text{dist}(x, y) \mid y - (0, t, 0) \in \mathbb{R}(0, 0, 1) + \mathbb{R}\mathbb{1}\}$. This is a 1-dimensional polyhedral complex, which is homeomorphic with $L(t)$ via orthogonal projection, consisting of five segments.

We will assume that our set S of sites is in general position and hence, in particular, each sweep line contains at most one site. A point $a = (a_1, a_2, a_3)$ is said to have been *visited* by the sweep line $L(t)$ if $a_2 - a_1 \leq t$.

The *beach line* $B(t)$ of S at time t is formed by the points (b_1, b_2, b_3) which lie on a parabola $P(s, t)$ for a visited point s such that $b_2 - b_1$ is maximal among all such points for a fixed value $b_3 - b_1$. That is, the beach line is formed by the right-most points on the parabolas spanned by the visited points and the sweep line; see Fig. 9. So $B(t)$ is a union of parabolic arcs; it is easy to see that each parabola contributes at most two arcs to the beach line at any time. Like a single parabola also the beach line $B(t)$ is homeomorphic to $L(t)$ via orthogonal projection. In the portion of $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ left to $B(t)$ the tropical Voronoi diagram of S is known at time t .

Observation 7 The beach line is a polygonal line with $O(n)$ segments.

The actual algorithm works as in the classical case. We maintain a priority queue of *site events* (when the sweep line visits a site) and *circle events* (when there is a candidate for a new vertex of the tropical Voronoi diagram). The total number of events is linear in n . As in the classical case, it is possible to relax the condition on general position by means of symbolic perturbation.

Theorem 8 *The beach line algorithm computes a tropical Voronoi diagram of n sites in $\mathbb{R}^3/\mathbb{R}\mathbb{1}$ in $O(n \log n)$ time and $O(n)$ space.*

Fig. 9 The beach line and the sweep line, for $t = 1$ (left) and $t = 3$ (right)

For the output we can choose, with the same complexity, between an abstract planar graph (encoding $\text{Vor}(S)$ topologically) and its piecewise linear embedding resulting from Theorem 6.

5.2 Polytrope Partitions

Let $S \subseteq \mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$ be a finite set of sites. From Theorem 6 we know that the tropical Voronoi diagram can be described in terms of (semi-)polytropes. For the definition and basic facts on polytropes, see Sect. 4.4 and [17]. The following takes inspiration from the *trapezoid map* data structure; see [7, §6].

Definition 2 A *polytrope partition* for S is a finite collection $\mathcal{C}$ of (perhaps unbounded) non-degenerate polytropes with disjoint interiors, covering $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$, such that:

1. each facet-defining hyperplane of any cell in $\mathcal{C}$ lies in the hyperplane arrangement $S + A_d$.
2. for each cell P in $\mathcal{C}$ and site $a \in S$ the restricted Voronoi region $\text{Vor}_S(a) \cap P$ is contained in a maximal cone $a + F$ of $a + \mathcal{F}(\mathbb{B}^d)$.

A *valid labeling* for $\mathcal{C}$ assigns to each cell $P \in \mathcal{C}$ a (partial) matching $\mathcal{L}_{\mathcal{C}}(P)$ of $S \times \mathcal{F}(\mathbb{B}^d)$ containing

$$\left\{ (a, F) \in S \times \mathcal{F}(\mathbb{B}^d) \mid (a + F) \cap P \cap \text{Vor}_S(a) \neq \emptyset \right\}.$$

The distance function $x \mapsto \text{dist}(x, a)$ to a fixed site a is piecewise linear, as it is linear in each translated cone $a + F$ for $F \in \mathcal{F}(\mathbb{B}^d)$. Polytrope partitions are designed to exploit this fact, so that the distances to the relevant sites in each cell are linear. In particular:

Observation 9 Let P be a cell in a polytrope partition $\mathcal{C}$ for S . Then for all $x \in P$ we have

$$\text{dist}(x, S) = \min_{a \in S} \lambda_{F_a}(x - a) \quad (16)$$

where λ_{F_a} is the linear function defined by restricting the distance to a on some maximal cone $a + F_a$ of $a + \mathcal{F}(\mathbb{B}^d)$ which contains $\text{Vor}_S(a) \cap P$. Thus computing

the restriction of $\text{Vor}(S)$ to the polytrope P amounts to finding the regions of linearity of the tropical polynomial $\min_{a \in S} \lambda_{F_a}(x)$. The latter can be obtained via an ordinary dual convex hull computation.

Note that the maximal cone F_a in the above is irrelevant if $\text{Vor}_S(a) \cap P = \emptyset$, and it is unique, by axiom (2), if $\text{Vor}_S(a) \cap P \neq \emptyset$. If the polytrope partition is equipped with a valid labeling, then this tells us the choice of the right cone F_a for each site, if it exists.

The following shows that polytrope partitions exist.

Example 6 The braid arrangement A_d consists of the $\binom{d+1}{2}$ ordinary hyperplanes $\{x \mid x_i = x_j\}$, where $i \neq j$. This gives rise to the *standard polytrope partition* $S + A_d$, which is finer than any other polytrope partition for S ; Fig. 8 shows an example for $d = 2$. This construction occurs in planar tropical convex hull algorithms; see [16, Figure 3].

And finally, the following lemma shows that valid labelings exist for every polytrope partition, if the sites S are in weak general position:

Lemma 11 *Let $\mathcal{C}$ be a polytrope partition for S . If S is in weak general position then there is a valid labeling of $\mathcal{C}$. Moreover, for d considered constant, a labeling of each polytropical cell has constant size, and it can be computed in $O(n)$ time.*

Proof Suppose that a valid labeling does not exist for some cell P . Either the set of pairs in (2) matches two cones with the same site, or matches two sites with the same cone of $\mathcal{F}(\mathbb{B}^d)$. The former cannot happen since P is full-dimensional and the cones in $\mathcal{F}(\mathbb{B}^d)$ only intersect in lower-dimensional polyhedral cones.

Then there are sites $a, b \in S$ and a maximal cone $F \in \mathcal{F}(\mathbb{B}^d)$ such that the sets $(a + F) \cap P \cap \text{Vor}_S(s)$ and $(b + F) \cap P \cap \text{Vor}_S(t)$ both are non-empty.

With the notation of (16) we have $\lambda_{F_a} = \lambda_{F_b}$; and we shortly write λ . Since the sites are in weak general position, we may assume that $\lambda(b) > \lambda(a)$. Picking $y \in (b + F) \cap P \cap \text{Vor}_S(b)$ yields

$$\text{dist}(y, b) \geq \lambda(b) > \lambda(a) \geq \text{dist}(y, S),$$

where the last inequality follows from (16). The resulting inequality $\text{dist}(y, b) > \text{dist}(y, S)$ implies that $y \notin \text{Vor}_S(b)$, which is a contradiction. Hence a valid labeling does exist.

To compute such labeling, we iterate through all the sites. For each site a , the candidate facet of $F_a \in \mathcal{F}(\mathbb{B}^d)$ is known by definition of the polytrope partition. To check if (a, F_a) is a labeling candidate, we need to determine if $(a + F_a) \cap P \cap \text{Vor}_S(a)$ is empty or not. This amounts to solving a linear program that has constant size (as d is a constant). It follows that the entire labeling can be computed in $O(n)$ time. $\square$

We aim at a first algorithm for computing a tropical Voronoi diagram in arbitrary dimension. This will employ the standard polytrope partition from Example 6.

Lemma 12 *If S is in weak general position and has size n then the trivial polytrope partition has*

$$(d+1)^{d-1}n^d + O(n^{d-1})$$

maximal cells, if we consider d a fixed constant.

Proof Pick a generic direction $v \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. The cells of the polytrope partition that are bounded in the direction of v are in a one-to-one correspondence with the vertices of the arrangement, by associating each polytrope with the optimum of the linear program maximizing $v^T x$. Since the number of vertices equals

$$(d+1)^{d-1}n^d - n \left((d+1)^{d-1} - 1 \right),$$

by Cayley's formula, it suffices to show that the number of unbounded cells is in $O(n^{d-1})$, as the bases of A_d correspond to spanning trees in the complete graph with d vertices. The unbounded cells intersect a hyperplane, H , normal to v that is far enough in the v direction. The cells intersecting H are the same as the cells in the restricted hyperplane arrangement, which is a $(d-1)$ -dimensional arrangement with $N = \binom{d+1}{2}n$ hyperplanes. The number of such cells is known to be in $O(N^{d-1})$, which agrees with $O(n^{d-1})$ as N depends linearly on n . $\square$

Remark 9 (Standard polytrope partition algorithm) This directly yields a first algorithm for computing a tropical Voronoi diagram of n sites in $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$ in $O(n^{d+1})$ time, as follows: First, we sort S along each of the $\binom{d+1}{2}$ directions $e_i + e_j$, in $O(n \log n)$ time. As in the Proof of Lemma 12 we pick a generic direction $v \in \mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. We can compute the vertices of the hyperplane arrangement $A_d + S$ in time $O(n^d)$ by enumerating all d -sets of independent directions, which can be derived from the oriented spanning tree of K_{d+1} , in constant time. For each of the d directions we choose an index $i \in [n]$.

Next we perturb each such vertex p by a small multiple of $-v$, and we collect the intersection of bands of contiguous parallel hyperplanes of $A_d + S$ that contain the perturbed point. This can be done in time $O(\log n)$ for each direction. In this way, we find those cells which are bounded in the direction of this particular v in linear output-dependent time. We repeat the same procedure for a set of directions $v_1, \dots, v_{d+1}$ which positively span the entire space $\mathbb{R}^{d+1}/\mathbb{R}\mathbb{1}$. Each polytrope cell will be bounded in at least one of these directions, and thus their enumeration is still in $O(n^d)$ for d fixed.

Then, for each polytrope P , we compute a corresponding labeling in time $O(n)$, by Lemma 11. Therefore, we can compute the standard polytrope partition, including labels, in time $O(n^{d+1})$ for fixed d . The tropical Voronoi diagram in each cell is an ordinary dual convex hull problem of constant size. This computation splits each polytrope in the partition into semi-polytropes. The convex hull problem can be solved in constant time, and hence this algorithm takes $O(n^{d+1})$ time, if d is considered a fixed constant.

We implemented this procedure in `polymake` [11], version 3.6.

Question 1 In the plane $\mathbb{R}^3/\mathbb{R}1$, we believe that ideas similar to the “trapezoidal maps” used in point location, see [7, §6.1], should yield polytrope partitions of linear size but we did not work out the details. More generally: Is there a polytrope partition of complexity better than $\Theta(n^d)$ in arbitrary dimension? One could hope for something in $O(n^{d/2})$, which is the worst-case complexity of Euclidean Voronoi diagrams.

5.3 An $O(n^d \log n)$ Randomized Incremental Algorithm in $\mathbb{R}^{d+1}/\mathbb{R}1$

We can improve the algorithm from Remark 9 by constructing a polytrope partition incrementally. The idea is to update an existing polytrope partition by including a new point and to employ randomization to improve the efficiency. Moreover, we will also produce a coarser polytrope partition than the standard one, but only by a constant in d .

A key ingredient is a new data structure that we call a *polytrope tree*. Throughout we assume that the set S of n sites forms a subset of $\mathbb{R}^{d+1}/\mathbb{R}1$ in general position. We fix the polytrope partition $\mathcal{C} := S + \mathcal{F}(\mathbb{B}^d)$, which is coarser than the standard polytrope partition but only by a factor which is constant in d ; see Example 6.

Definition 3 A *polytrope tree* for S is a (rooted) tree T such that

1. for each leaf ℓ there is a polytrope cell $P(\ell)$ of $\mathcal{C}$;
2. for each interior node i there is a site $a(i) \in S$ and a polytrope $P(i)$.

These satisfy the following consistency conditions:

- for the root node r of T we have $P(r) = \mathbb{R}^{d+1}/\mathbb{R}1$, which may be seen as an unbounded polytrope;
- the map $\ell \mapsto P(\ell)$ is a bijection between the set of leaves of T and the set of polytropes in $\mathcal{C}$;
- the map $i \mapsto a(i)$ is a surjection from the interior nodes onto the set S ;
- if i is an interior node with children $c_1, \dots, c_k$, then $P(c_1), \dots, P(c_k)$ form the maximal cells of $(a(i) + \mathcal{F}(\mathbb{B}^d)) \cap P(i)$.

It is easy to construct a polytrope tree for S , and its purpose is to speed up the computation of a valid labeling. This will reduce the algorithmic complexity from $O(n^{d+1})$ to $O(n^d \log n)$.

For the incremental update to insert a new site $b \notin S$ we maintain a stack Σ of unvisited nodes in a given polytrope tree for Σ and process it as follows:

- the stack Σ is initialized with the root node r ;
- we remove the top node q from the stack Σ unless it is empty;
- if q is an interior node such that $P(q)$ intersects more than one maximal cone of $s + \mathcal{F}(\mathbb{B}^d)$, then we push the children of q onto the stack Σ ;
- if q is a leaf such that $P(\ell)$ intersects more than one maximal cone of $p + \mathcal{F}(\mathbb{B}^d)$, then we create the intersections of $P(\ell)$ with $s + \mathcal{F}(\mathbb{B}^d)$ as new leaves, which now become children of q , and we set $a(q) \leftarrow b$.

Note that an interior node q with $P(q)$ contained in a unique maximal cone of $a + \mathcal{F}(\mathbb{B}^d)$ is kept unchanged, and its children will not be visited. The following is the essential part of the complexity analysis.

Proposition 7 *Let T be a polytrope tree created in the way explained above, where the n sites in S are processed in uniformly random order. Then the expected height of T is of order $O(\log n)$, if d is considered a fixed constant.*

Proof Let P be a polytrope in the polytrope partition $\mathcal{C}(S)$. For each ordering $\pi : [n] \rightarrow S$ of S we have a polytrope tree $T(S, \pi)$ with P as a leaf. By induction on n we will show:

$$E[h_{T(S, \pi)}(P)] \leq d(d+1) \sum_{i=1}^n \frac{1}{i} \in O(\log n), \quad (17)$$

where the expectation $E[\cdot]$ is taken uniformly over all $n!$ orderings of S , and $h_{T(S, \pi)}(P)$ is the depth of the leaf of P in $T(S, \pi)$.

We proceed by backwards analysis. Let $S' \subset S$ be the subset of sites that lie in some facet-defining hyperplane of P . Since P has at most $d(d+1)$ facets and (by general position) their corresponding hyperplanes contain each exactly one point of S , we have $|S'| \leq d(d+1)$. Thus, the probability that the height $h_T(P)$ increases in the last insertion is at most $d(d+1)/n$. Since the increase is by exactly one, we have

$$E[h_{T(S, \pi)}(P)] \leq E[h_{T(S \setminus \pi(n), \pi_{[n-1]})}(P')] + \frac{d(d+1)}{n},$$

where P' is the polytrope containing P in the polytrope partition before the last insertion. By induction hypothesis

$$E[h_{T(S \setminus \pi(n), \pi_{[n-1]})}(P')] \leq d(d+1) \sum_{i=1}^{n-1} \frac{1}{i}.$$

The last two formulas give Eq. (17). $\square$

Corollary 9 *The above method constructs a polytrope tree for the polytrope partition $S + \mathcal{F}(\mathbb{B}^d)$ in expected time $O(n^d \log n)$ and space $O(n^d)$, for d constant.*

Proof The algorithm that inserts a new site a into the tree only visits nodes that are above some leaf requiring an update. For each such leaf ℓ the polytrope $P(\ell)$ intersects one of the $d(d+1)$ hyperplanes in $a + A_d$. This implies that there are $O(n^{d-1})$ of them. Since the expected depth of every leaf is $O(\log n)$ it requires expected time $O(n^{d-1} \log n)$ for inserting a . Hence the total complexity for n sites amounts to $O(n^d \log n)$. $\square$

In order to compute the tropical Voronoi diagram, we also need to compute the labeling of this polytrope partition. The naive way is to compute the labeling for each leaf as we did in Remark 9.

A slight improvement is to compute the labeling during the depth-first-search (DFS) exploration of the tree at each insertion of a new site. But in this way, even if an interior node is completely contained in only one cone of the fan $a + \mathcal{F}(\mathbb{B}^d)$, we need to descend

to its subtree in order to update the labels. This would slow the algorithm down to $\Theta(n^{d+1})$ because each insertion will have to iterate through all the leaves.

A better way is to compute the labeling lazily. To this end we equip each interior node i with a partial labeling $\mathcal{L}_C(i)$. With each new insertion, we proceed as we just explained, but we do not cascade down the label updates. Only once all sites in S have been inserted we cascade the partial labels, updating them in DFS order. This takes $O(n^d \log n)$ time to compute the polytropes and the lazy labelings, plus $O(n^d)$ time to cascade the lazy labelings down in the tree, for a total time complexity of $O(n^d \log n)$ time. This gives our final result.

Theorem 10 *There is a randomized incremental algorithm for computing tropical Voronoi diagrams of n sites in $\mathbb{R}^{d+1}/\mathbb{R}\mathbf{1}$ in general position with expected time complexity $O(n^d \log n)$ and space complexity $O(n^d)$, for d constant.*

Acknowledgements We thank Günter Rote and three anonymous referees for useful comments on a first version of this paper. Especially, one referee brought the reference [3] to our attention.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Daniele Alessandrini, *Logarithmic limit sets of real semi-algebraic sets*, Adv. Geom. **13** (2013), no. 1, 155–190. MR 3011539.
2. Xavier Allamigeon, Pascal Benchimol, Stéphane Gaubert, and Michael Joswig, *Log-barrier interior point methods are not strongly polynomial*, SIAM J. Appl. Algebra Geom. **2** (2018), no. 1, 140–178.
3. Omid Amini and Madhusudan Manjunath, *Riemann-Roch for sub-lattices of the root lattice A_n* , Electron. J. Combin. **17** (2010), no. 1, Research Paper 124, 50. MR 2729373.
4. Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee, *Voronoi diagrams and Delaunay triangulations*, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2013. MR 3186045.
5. Matthew Baker and Serguei Norine, *Riemann-Roch and Abel-Jacobi theory on a finite graph*, Adv. Math. **215** (2007), no. 2, 766–788. MR 2355607.
6. L. Paul Chew and Robert L. Scot Dyrdsdale III, *Voronoi diagrams based on convex distance functions*, Proceedings of the first annual symposium on Computational geometry, ACM, 1985, pp. 235–244.
7. Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars, *Computational geometry*, third ed., Springer-Verlag, Berlin, 2008, Algorithms and applications. MR 2723879.
8. Jules Depersin, Stéphane Gaubert, and Michael Joswig, *A tropical isoperimetric inequality*, Sémin. Lothar. Combin. – FPSAC 2017 – **78B** (2017), Art. 27, 12. MR 3678609.
9. Mike Develin, *The moduli space of n tropically collinear points in $\mathbb{R}^d$* , Collect. Math. **56** (2005), no. 1, 1–19. MR 2131129.
10. Steven Fortune, *A sweepline algorithm for Voronoi diagrams*, Algorithmica **2** (1987), no. 2, 153–174. MR895442 (88e:68101).
11. Ewgenij Gawrilow and Michael Joswig, *polymake: a framework for analyzing convex polytopes*, Polytopes—combinatorics and computation (Oberwolfach, 1997), DMV Sem., vol. 29, Birkhäuser, Basel, 2000, pp. 43–73. MR1785292 (2001f:52033).

12. Chan He, Horst Martini, and Senlin Wu, *On bisectors for convex distance functions*, Extracta Math. **28** (2013), no. 1, 57–76. MR 3135681.
13. Christian Icking, Rolf Klein, Ngoc-Minh Lê, and Lihong Ma, *Convex distance functions in 3-space are different*, Fund. Inform. **22** (1995), no. 4, 331–352, Computational geometry (San Diego, CA, 1993). MR 1360951.
14. Christian Icking, Rolf Klein, Ngoc-Minh Lê, Lihong Ma, and Francisco Santos, *On bisectors for convex distance functions in 3-space*, CCCG, (1999).
15. Philipp Jell, Claus Scheiderer, and Josephine Yu, *Real tropicalization and analytification of semialgebraic sets*, Int. Math. Res. Not. IMRN (2020), Published online: 29 May 2020.
16. Michael Joswig, *Tropical halfspaces*, Combinatorial and computational geometry, Math. Sci. Res. Inst. Publ., vol. 52, Cambridge Univ. Press, Cambridge, 2005, pp. 409–431. MR2178330 (2006g:52012).
17. Michael Joswig and Katja Kulas, *Tropical and ordinary convexity combined*, Adv. Geometry **10** (2010), 333–352.
18. Michael Joswig and Georg Loho, *Weighted digraphs and tropical cones*, Linear Algebra Appl. **501** (2016), 304–343.
19. J. J. Kosowsky and Alan L. Yuille, *The invisible hand algorithm: Solving the assignment problem with statistical physics*, Neural Networks **7** (1994), no. 3, 477–490.
20. Thomas Lam and Alexander Postnikov, *Alcoved polytopes. I*, Discrete Comput. Geom. **38** (2007), no. 3, 453–478. MR2352704.
21. Anthea Monod, Bo Lin, Ruriko Yoshida, Qiwen Kang, *Tropical geometry of phylogenetic tree space: a statistical perspective*, (2020), Preprint [arXiv:1805.12400v6](https://arxiv.org/abs/1805.12400v6)
22. Diane Maclagan and Bernd Sturmfels, *Introduction to tropical geometry*, Graduate Studies in Mathematics, vol. 161, American Mathematical Society, Providence, RI, 2015. 3287221
23. Horst Martini and Konrad J Swanepoel, *The geometry of Minkowski spaces—a survey. part ii*, Expositiones mathematicae **22** (2004), no. 2, 93–144.
24. Grigory Mikhalkin, *Enumerative tropical algebraic geometry in $\mathbb{R}^2$* , J. Amer. Math. Soc. **18** (2005), no. 2, 313–377. 2137980 (2006b:14097).
25. David Speyer and Bernd Sturmfels, *The tropical Grassmannian*, Adv. Geom. **4** (2004), no. 3, 389–411. MR2071813 (2005d:14089)
26. Rade T. Živaljević, *Topological methods in discrete geometry*, Handbook of Discrete and Computational Geometry (Csaba D. Tóth, Jacob E. Goodman, and Joseph O’Rourke, eds.), CRC Press, Boston, 3rd edition.
27. Peter Widmayer, Ying-Fung Wu, and Chak-Kuen Wong, *On some distance problems in fixed orientations*, SIAM Journal on Computing **16** (1987), no. 4, 728–746.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.