Maximal f-Vectors of Minkowski Sums of Large Numbers of Polytopes

Christophe Weibel

Abstract It is known that in the Minkowski sum of r polytopes in dimension d, with $r < d$, the number of vertices of the sum can be as high as the product of the number of vertices in each summand. However, the number of vertices for sums of more polytopes was unknown so far.

In this paper, we study sums of polytopes and prove a linear relation between the number of faces of a sum of r polytopes in dimension d, with $r \geq d$, and the number of faces in the sums of less than d of the summand polytopes. We deduce from this result a bound on the maximum possible number of vertices of the Minkowski sum of any number of polytopes in any dimension. In particular, the linear relation implies that a sum of r polytopes in dimension d, where summands have n vertices in total, has less than $\binom{n}{d-1}$ vertices, even when $r \geq d$.

Finally, we present a construction for any given number of vertices in summands and show that no other sum can achieve more vertices, establishing a precise tight bound.

1 Introduction

The *Minkowski sum* of two polytopes, P_1 and P_2, is defined as $\{x_1 + x_2 : x_1 \in P_1, x_2 \in P_2\}$. Minkowski sums are of interest in various fields of theoretical and applied mathematics. While some applications only require sums of two polytopes in low dimensions (e.g., motion planning [5, 8]), others use iterative sums of many polytopes

---

C. Weibel (✉)
Computer Science, Dartmouth College, HB 6211, Hanover, NH 03755, USA
e-mail: weibel@cs.dartmouth.edu

C. Weibel
e-mail: christophe.weibel@gmail.com

in higher dimensions [9, 12]. It is therefore desirable to study the complexity of such sums.

A trivial bound on the number of vertices of a sum is found as follows. Every vertex of a Minkowski sum decomposes into a sum of vertices of the summands. Therefore, there cannot be more vertices in the sum than there are possible decompositions. Thus, the *trivial bound* on the number of vertices of a Minkowski sum is the product of the number of vertices of the summands. That is, if $P_1, \ldots, P_r$ are polytopes, and $f_0(P)$ is the number of vertices of a polytope P, then $f_0(P_1 + \cdots + P_r) \leq f_0(P_1) \times \cdots \times f_0(P_r)$.

If we sum r polytopes in dimension d, with $r < d$, then the trivial bound is tight. That is, it is possible to choose summands with any number of vertices so that their sum has as many vertices as the trivial bound $f_0(P_1) \times \cdots \times f_0(P_r)$ [2].

However, if we sum r polytopes in dimension d with $r \geq d$, the trivial bound cannot be reached, except when summing d segments [10, 11]. We assume here as in the rest of the article that polytopes have at least two vertices, since summands of only one vertex can be ignored without changing the properties of the sum.

The f-vector of a polytope encodes the number of the polytope faces of different dimensions. For fixed f-vectors of summands, maximum f-vectors in the sum can always be reached by summands in general orientations [3], because slightly perturbing summands so that they are in general orientations can only increase the number of faces in the sum. Thus, upper bounds on the number of faces of sums of polytopes in general orientations also apply to sums of polytopes that are not in general orientations. We assume hereafter that summands are in general orientations and full-dimensional.

We recently presented in [1] a tight bound on the maximum number of vertices and facets in sums of three-dimensional polytopes, by showing that a bound on the number of vertices of a sum of r summands can be deduced from the number of vertices of the summands and the number of vertices of sums of each of the $\binom{r}{2}$ pairs of summands.

The basic reasoning of this previous result is to define a unique witness, called *western-most corner*, for all but two vertices of a polytope. These witnesses have the property that a western-most corner for a Minkowski sum of any number of summands is also necessarily a western-most corner for the sum of some pair of summands. So we can find the number of western-most corners of the total sum, and its number of vertices, by examining individual summands and sums of two summands only.

In this paper we extend this reasoning. Our main result, presented in Theorem 1, considerably generalizes the previous one, as it applies to all faces and any dimension. It is a linear relation between the number of faces of a sum of r polytopes and the number of faces in the sums of less than d of the summand polytopes:

Theorem 1 *Let $P_1, \ldots, P_r$ be d-dimensional polytopes in general orientations, $r \geq d$, and each polytope full-dimensional. For any k in $0, \ldots, d-1$,*

$$f_k(P_1 + \cdots + P_r) - \alpha = \sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{r-1-j}{d-1-j} \sum_{S \in \mathcal{C}_j^r} \big(f_k(P_S) - \alpha\big),$$

where $f_k(P)$ indicates the number of k-dimensional faces of a polytope P, $\mathcal{C}_j^r$ is the family of subsets of $\{1, \ldots, r\}$ of cardinality j, P_S is the sum of polytopes $\sum_{i \in S} P_i$; $\alpha = 2$ if $k = 0$ and d is odd, $\alpha = 0$ otherwise.

A slightly more general result also applies when summands are not full-dimensional.

The intuitive explanation of the theorem is based on the inclusion–exclusion principle. For any face of the whole sum, we can find a witness of its existence by examining the faces of the same dimension in sums of $d - 1$ summands. However, if that witness exists in some sum of $d - 2$ summands, we will find it in many different sums of $d - 1$ summands. So we need to offset this by removing an appropriate number of times the witnesses in sums of $d - 2$ summands. But that in turn removes too many times witnesses that exist in some sum of $d - 3$ summands, so we need to add them back, etc. As any face of the whole sum does have at least one witness in sums of $d - 1$ summands, the alternate sum is smaller than its term for $j = d - 1$:

Corollary 2 *Let $P_1, \ldots, P_r$ be d-dimensional polytopes, $r \geq d$, and each polytope full-dimensional. For any k in $0, \ldots, d - 1$,*

$$f_k(P_1 + \cdots + P_r) \leq \sum_{S \in \mathcal{C}_{d-1}^r} f_k(P_S).$$

We deduce from this result upper bounds on the maximum possible number of vertices in Minkowski sums. We find in particular that a sum of r polytopes in dimension d, $r \geq d$, where summands have n vertices in total, has less than $\binom{n}{d-1}$ vertices. In the case where each summand has at most n vertices, the number of vertices of the sum is less than $\binom{r}{d-1} n^{d-1}$. The previously known bound for vertices of Minkowski sums was in $O(r^{d-1} n^{2(d-1)})$ [4]. This has important implications on the usefulness of methods that use sums of many polytopes in comparatively low dimensions [9, 12], since we obtain a better bound on their complexity. The actual complexity of $O(r^{d-1} n^{d-1})$ could also be deduced from Theorem 3.3 of [7], but Corollary 2 is more precise, since it provides the constant for the order of complexity.

We further refine this upper bound by proving the *exact* maximum possible number of vertices of a Minkowski sum of full-dimensional polytopes for fixed number of vertices of summands:

Theorem 3 *Let $P_1, \ldots, P_r$ be d-dimensional polytopes, $r \geq d$, and each polytope full-dimensional. Then*

$$f_0(P_1 + \cdots + P_r) \leq \alpha + \sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{r-1-j}{d-1-j} \sum_{S \in \mathcal{C}_j^r} \left(\prod_{i \in S} f_0(P_i) - \alpha \right),$$

where $\alpha = 2$ if d is odd, and $\alpha = 0$ if d is even. This bound is tight.

We achieve this result by giving a construction that reaches this bound and proving that no sum of full-dimensional polytopes can have more vertices.

The rest of the article is organized as follows. We briefly present the background theory in Sect. 2. In Sect. 3, we introduce the concepts of west and western-most corner in three dimensions and then formally extend them to higher dimensions in Sect. 4. We examine in Sect. 5 the maximum possible number of vertices of a Minkowski sum. Most of the technical proofs are deferred to Sect. 6. Finally, we summarize our results in Sect. 7.

2 Minkowski Sums

Let $P_1, \ldots, P_r$ be r polytopes. Their Minkowski sum is the polytope defined as $P_1 + \cdots + P_r = \{x_1 + \cdots + x_r : x_i \in P_i \; \forall i\}$. We assume hereafter that every polytope is full-dimensional.

A nontrivial *face* of a polytope P in dimension d is the intersection of P with a support hyperplane of P. Vertices, edges, facets, and ridges are the faces of dimension $0, 1, d-1$, and $d-2$, respectively. Thus, we can associate with each vector l in the unit sphere S^{d-1} a face of the polytope, which is the intersection of the polytope with the support hyperplane to which the vector is outwardly normal: $\mathcal{S}(P; l) = \{x \in P : \langle l, x \rangle \geq \langle l, y \rangle \; \forall y \in P\}$, where $\langle \cdot, \cdot \rangle$ represents the standard Euclidean inner product.

Conversely, each face F of a d-dimensional polytope P can be associated with a region of the sphere S^{d-1}, called the *normal region*, which is the set of unit vectors outwardly normal to some support hyperplane of P, whose intersection with P is F: $\mathcal{N}(F; P) = \{l \in S^{d-1} \mid F = \mathcal{S}(P; l)\} = \{l \in S^{d-1} \mid \langle l, x \rangle > \langle l, y \rangle \; \forall x \in F, \; y \in P \setminus F\}$. Alternatively, the normal region of a face can be defined as the intersection of its relatively open normal cone with S^{d-1}. The normal region of a face of dimension k is a relatively open subset of S^{d-1} of dimension $d-1-k$. In particular, the normal region of a facet of P is a single point of S^{d-1}, corresponding to the unit vector outwardly normal to the facet.

We call a subset of the sphere S^{d-1} *spherically convex* if for any two points in the subset, any shortest arc of great circle between the two points is inside the subset.[1] If the polytope P is full-dimensional, the normal regions of faces of P are all disjoint, relatively open and spherically convex. They determine a subdivision of S^{d-1} into a spherical cell complex, which we call the *Gaussian map* of the polytope: $\mathcal{G}(P) = \{\mathcal{N}(F; P) : F \text{ face of } P\}$.

A property of Minkowski sums is that faces of the sum have a unique *decomposition* into faces of the summand. Let F be a face of the Minkowski sum $P = P_1 + \cdots + P_r$, and l be in $\mathcal{N}(F; P)$. Then $F = F_1 + \cdots + F_r$, where $F_i = \mathcal{S}(P_i; l)$ is a face of P_i. We deduce that the normal region of a face of the sum is equal to the intersection of the normal regions of faces in its decomposition: $\mathcal{N}(F; P) = \mathcal{N}(F_1; P_1) \cap \cdots \cap \mathcal{N}(F_r; P_r)$. Thus the Gaussian map of the sum is the *common refinement* of the Gaussian map of the summands:

$$\mathcal{G}(P_1 + \cdots + P_r) = \{\mathcal{N}(F_1; P_1) \cap \cdots \cap \mathcal{N}(F_r; P_r) : F_i \text{ face of } P_i\}.$$

---

[1] There exist different definitions of convexity on a sphere. Note that according to this one, the only convex set containing antipodal points is the whole sphere.

We can study the number of faces of a polytope by studying the number of cells of its Gaussian map, as the polytope and its Gaussian map are dual structures.

We say that a face of a Minkowski sum has an *exact decomposition* $F = F_1 + \cdots + F_r$ when its dimension is the sum of the dimension of the faces in its decomposition: $\dim(F) = \dim(F_1) + \cdots + \dim(F_r)$. That is, the decomposition is exact when there are no two parallel segments inside different faces in the decomposition. We say that polytopes are *in general orientations* when all faces of their Minkowski sum have an exact decomposition.

For fixed f-vectors of summands, the maximum number of faces of any dimension in the sum can always be reached by summands in general orientations [3]. Thus, bounds on the number of faces of such sums apply to all other sums.

Let F be a face of the Minkowski sum $P = P_1 + \cdots + P_r$ of d-dimensional polytopes in general orientations. The face F decomposes into a sum $F_1 + \cdots + F_r$ of faces of the summands, with $\dim(F) = \dim(F_1) + \cdots + \dim(F_r)$. Even if $r \geq d$, there are at most $\dim(F)$ faces in the decomposition that are of a dimension more than 0. Let the *support* $I_F \subseteq \{1, \ldots, r\}$ of F be the set of indices of these faces, with $|I_F| \leq \dim(F)$. Note that for any subface G of F, G decomposes into a sum $G_1 + \cdots + G_r$, where $G_i \subseteq F_i$ for all i; and so, $I_G \subseteq I_F$.

For any subset $S = \{i_1, \ldots, i_s\}$ of $\{1, \ldots, r\}$, let us define the *partial sum* $P_S = P_{i_1} + \cdots + P_{i_s}$.

Lemma 4 *Let F be a facet of a Minkowski sum $P = P_1 + \cdots + P_r$ of d-dimensional polytopes in general orientations. Its normal region $\mathcal{N}(F; P)$ is a node of $\mathcal{G}(P)$. Then, $\mathcal{N}(F; P)$ is also a node of the Gaussian map $\mathcal{G}(P_S)$ of a partial sum if and only if $I_F \subseteq S$.*

Proof Let $F_1 + \cdots + F_r$ be the decomposition of F, with $\dim(F_i) > 0$ if and only if $i \in I_F$. Since the summands are in general orientations, the decomposition is exact, and $\dim(F) = d - 1 = \dim(F_1) + \cdots + \dim(F_r) = \sum_{i \in I_F} \dim(F_i)$. The normal region $\mathcal{N}(F; P)$ contains a single unit vector l, and $\mathcal{N}(F; P)$ is a node of $\mathcal{G}(P_S)$ if and only if $\dim(\mathcal{S}(P_S; l)) = d - 1$. Again, the decomposition is exact, and $\dim(\mathcal{S}(P_S; l)) = \sum_{i \in S} \dim(\mathcal{S}(P_i; l)) = \sum_{i \in S} \dim(F_i)$. Since $\dim(F_i) > 0$ if and only if $i \in I_F$, and $\sum_{i \in I_F} \dim(F_i) = d - 1$, the result follows. □

3 Sums of Polytopes in Dimension 3

We informally present in this section the argument for the proof of Theorem 1 for three dimensions, as it is easier to grasp and facilitates the comprehension of the proof for arbitrary dimensions, which we present in Sect. 4. The result for three dimensions has already been published [1].

In dimension 3, the Gaussian map of a polytope is a spherical cell complex of S^2, which can be described as a planar graph embedded in S^2. The normal regions of facets, edges, and vertices of the polytope are nodes, edges, and faces of the graph, respectively. Note that the normal regions of edges of the polytope, edges of the graph, are arcs of great circles of S^2. The Gaussian map of a Minkowski sum is the *overlay* of the Gaussian maps of the summands.
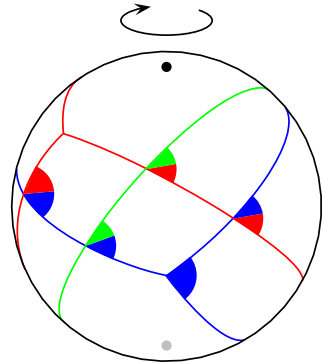
Let $P = P_1 + \cdots + P_r$ be a sum of three-dimensional polytopes in general orientations. We choose on S^2 two antipodal points that do not belong to any great circle containing an edge of $\mathcal{G}(P)$. In particular, the points are in the interior of two distinct faces of $\mathcal{G}(P)$. We call these two points *north pole* and *south pole*. We define *west* in the usual way with respect to the poles as a direction turning about an axis going through the poles, clockwise from the north pole.

For any spherically convex subset C of S^2 that does not contain either pole, we define the *western-most point* of C as the point in the closure of C that is to the west of all points in C. We also define the *western-most corner* of C. If C is not a vertex, its western-most corner is the subset of C at distance less than ε of its western-most point, where $\varepsilon > 0$ is smaller than the distance between any two nonincident cells in $\mathcal{G}(P)$. If C is a vertex, its western-most corner is its western-most point, that is, the vertex itself. Note that if C is a cell of $\mathcal{G}(P)$, the western-most point is a node of $\mathcal{G}(P)$ incident to C; it is unique, because otherwise an edge of $\mathcal{G}(P)$ would be contained in a great circle going through the poles.

The normal region $\mathcal{N}(F; P)$ of a facet F of the Minkowski sum P is a node of $\mathcal{G}(P)$. Because the summands are in general orientations, there are no three edges of different $\mathcal{G}(P_i)$ intersecting in a single point. Thus, a node in $\mathcal{G}(P)$ is either a node in some $\mathcal{G}(P_i)$, in which case $I_F = \{i\}$, or it is the intersection of two edges in some $\mathcal{G}(P_i + P_j)$, in which case $I_F = \{i, j\}$ (see Fig. 1). So a western-most corner in $\mathcal{G}(P)$ is always a western-most corner in some $\mathcal{G}(P_i)$, or a western-most corner in some $\mathcal{G}(P_i + P_j)$, whose western-most point is the intersection of two edges. Thus, we can find the number of western-most corners of $\mathcal{G}(P)$ by counting those in all $\mathcal{G}(P_i)$ and those in all $\mathcal{G}(P_i + P_j)$ whose western-most point is an intersection of edges. But then, the western-most corners in $\mathcal{G}(P_i + P_j)$ also include corners whose western-most point is a node of $\mathcal{G}(P_i)$ or $\mathcal{G}(P_j)$. Denoting as $w_k(g)$ the number of western-most corners of k-dimensional cells in a Gaussian map g, this means that

$$w_k\big(\mathcal{G}(P)\big) = \sum_{i=1}^{r} w_k\big(\mathcal{G}(P_i)\big) + \sum_{1 \leq i < j \leq r} \big(w_k\big(\mathcal{G}(P_i + P_j)\big) - w_k\big(\mathcal{G}(P_i)\big) - w_k\big(\mathcal{G}(P_j)\big)\big)$$

$$= \sum_{1 \leq i < j \leq r} w_k\big(\mathcal{G}(P_i + P_j)\big) - (r-2) \sum_{i=1}^{r} w_k\big(\mathcal{G}(P_i)\big), \quad k = 0, 1, 2.$$

In words, we sum the number of western-most corners of different $\mathcal{G}(P_i + P_j)$ and subtract $(r - 2)$ times the western-most corners in the $\mathcal{G}(P_i)$, since they are each counted $(r - 1)$ times in the first sum.

There is one distinct western-most corner for every cell of a Gaussian map except the two two-dimensional cells that contain a pole. As the cells of a Gaussian map correspond to faces of the underlying polytope, we have that for any polytope $\mathcal{P}$, $w_0(\mathcal{G}(\mathcal{P})) = f_2(\mathcal{P})$, $w_1(\mathcal{G}(\mathcal{P})) = f_1(\mathcal{P})$, and $w_2(\mathcal{G}(\mathcal{P})) = f_0(\mathcal{P}) - 2$. Replacing w_k in the above equation, we get Theorem 1 for $d = 3$.

We emphasize a subtle detail of the argument. Let us say that a point p in the closure of a subset C of the sphere S^2 is a *local optimum* of C if p is the western-most point for the intersection of C with some open set containing p. The reason we use the direction west is that the level curves for west, the meridians of geography, are arcs of great circles, i.e., geodesics; they intersect only once any other geodesic inside a spherically convex set. This guarantees that all local optima are also western-most points. This would not be the case had we used the direction south, because the level curves for south are not geodesics. Using west guarantees that a node that is a western-most point for $\mathcal{G}(P)$ is still a western-most point when it exists in some map $\mathcal{G}(P_i)$ or $\mathcal{G}(P_i + P_j)$.

4 Generalization to Higher Dimensions

In this section, we generalize the argument of Sect. 3 to higher dimensions. First, we extend the definition of west and western-most corners. We prove that the extension has the same property that (about) every cell of a Gaussian map has a single western-most corner. We also prove that the western-most corner of some cell in the Gaussian map of a Minkowski sum is also a western-most corner of some cell in any Gaussian map of a partial sum, where its western-most point is a node of the map. Finally, we present the formula that allows us to count the number of western-most corners of the Gaussian map of the Minkowski sum.

Proof of Theorem 1 We first present a sketch of the proof. Let $P = P_1 + \cdots + P_r$ be a Minkowski sum of d-dimensional polytopes in general orientations. Because the summands are in general orientations, a node of $\mathcal{G}(P)$ is also a node of $\mathcal{G}(P_S)$ if and only if the support of its underlying facet is contained in S (Lemma 4). This implies that in all $\mathcal{G}(P_S)$ where a node exists, the local geometry of the map around the node is the same as in $\mathcal{G}(P)$. Therefore, if the node is a local optimum in $\mathcal{G}(P)$, it is also a local optimum in any $\mathcal{G}(P_S)$ of which it is a node. This implies that a western-most corner in $\mathcal{G}(P)$ is also a western-most corner in $\mathcal{G}(P_S)$ if and only if its western-most point in $\mathcal{G}(P)$ exists in $\mathcal{G}(P_S)$, i.e., if and only if S contains the support of the underlying facet of the node. This is what allows us to use a counting argument to deduce the number of western-most corners of $\mathcal{G}(P)$ from the number of western-most corners of the Gaussian map of partial sums. Since there is one western-most corner per cell of the Gaussian map and face of the underlying polytope, this allows us to find the number of faces of the sum P.

We now extend the definition of west. The Gaussian map $\mathcal{G}(P)$ subdivides the sphere S^{d-1} into a spherical cell complex. The normal regions of ridges of P are

arcs of great circles on S^{d-1}, edges of the Gaussian map $\mathcal{G}(P)$. Each of these arcs of great circles is contained in the two-dimensional subspace that is orthogonal to the underlying ridge. In the d-dimensional space containing S^{d-1}, we choose a linear subspace U of dimension $d-2$, so that its intersection with every two-dimensional subspace containing an edge of $\mathcal{G}(P)$ is just the origin. The next lemma shows that this is always possible.

Lemma 5 *In a d-dimensional space, for any finite family $\{U_1, \ldots, U_n\}$ of two-dimensional linear subspaces, it is possible to find a linear subspace U of dimension $d-2$, such that $U \cap U_i = \{\mathbf{0}\}$ for any i in $1, \ldots, n$.*

Proof The orthogonal complements $U_i^{\perp}$ of the subspaces U_i in the family are of dimension $d-2$. If we choose a vector $\mathbf{u}$ that is not in these orthogonal complements, then for any i, $\mathrm{span}(\{\mathbf{u}\} \cup U_i^{\perp})$ is of dimension $d-1$. If we choose now a vector $\mathbf{v}$ that is not in any of these subspaces of dimension $d-1$, then for any i, $\mathrm{span}(\{\mathbf{u}, \mathbf{v}\} \cup U_i^{\perp})$ is the d-dimensional space. Define $U = \mathrm{span}(\{\mathbf{u}, \mathbf{v}\})^{\perp}$. If a vector is in $U \cap U_i$, it is orthogonal to $\mathrm{span}(\{\mathbf{u}, \mathbf{v}\} \cup U_i^{\perp})$, which is the whole space, and so it is the origin $\mathbf{0}$. □

This lemma is a simple extension of the fact that in three dimensions, for any number of planes going through the origin, we can choose a vector that is in none of the planes.

Note that the intersection of U with S^{d-1} is affinely isomorphic to the sphere S^{d-3}. If $d = 3$, the intersection is the two antipodal points on S^2 that we named north and south poles in Sect. 3, and if $d = 2$, it is the empty set because $U = \{\mathbf{0}\}$.

We choose an orthonormal basis $\mathbf{e_1}, \ldots, \mathbf{e_d}$ of the d-dimensional space, such that $U = \mathrm{span}(\{\mathbf{e_3}, \ldots, \mathbf{e_d}\})$. We then define successive polar parameterizations of the spheres S^k, $1 \le k \le d-1$ as follows:

$$S^1 = \{\sin(\theta_1)\mathbf{e_1} + \cos(\theta_1)\mathbf{e_2} : \theta_1 \in [0, 2\pi)\},$$

$$S^k = \{\sin(\theta_k)S^{k-1} + \cos(\theta_k)\mathbf{e_{k+1}} : \theta_k \in [0, \pi]\}, \quad k = 2, \ldots, d-1.$$

Note that a point of S^{d-1} is in U if and only if $\sin(\theta_j) = 0$ for some $j = 2, \ldots, d-1$. For any point of S^{d-1} not in U, we define the direction west as $\dot{\theta}_1$, the direction of augmentation of θ_1. Note that for $d = 3$, it is equivalent to the definition of Sect. 3, and for $d = 2$, it is a direction running around S^1. In S^{d-1}, west is not defined on the subspace U because $\dot{\theta}_1 = 0$, so that the intersection of U with S^{d-1} is a sphere of dimension $d-3$ that plays the same role as poles in three dimensions.

Recall that in our parameterization of S^{d-1}, θ_1 is in $[0, 2\pi)$. Formally, for any points p and q of S^{d-1} that are not in U, we say that p is *to the west* of q if $\theta_1(p) \in [\theta_1(q), \theta_1(q) + \pi]$ and $\theta_1(q) < \pi$, or if $\theta_1(p) \in [\theta_1(q), 2\pi) \cup [0, \theta_1(q) - \pi]$ and $\theta_1(q) \ge \pi$.

For any spherically convex subset C of S^{d-1} that does not intersect U, we define the *western-most point* of C as the point in the closure of C that is to the west of all points in C. The next lemma, proved in Sect. 6.1, shows that the western-most point exists.

Lemma 6 *If a spherically convex subset C of S^{d-1} does not intersect U, it is contained in a hemisphere defined by $\theta_1 \in [\alpha, \alpha + \pi]$ or $\theta_1 \in [0, \alpha] \cup [\alpha + \pi, 2\pi)$ for some $\alpha \in [0, \pi)$.*

We also define as *western-most corner* of C the subset of C at distance less than ε of the western-most point, where $\varepsilon > 0$ is smaller than the distance between any two nonincident cells in $\mathcal{G}(P)$. Note that the western-most point of C is also the western-most point of the western-most corner of C.

Recall that the Gaussian map of a polytope is a subdivision of S^{d-1} into a spherical cell complex. For any cell C of $\mathcal{G}(P)$ that does not intersect U, the western-most point of C is a node incident to C, which is unique because otherwise there would be a great circle containing an edge of $\mathcal{G}(P)$ and intersecting U, which contradicts the way we chose U. As a consequence, there is one unique western-most corner for each cell of a Gaussian map that does not intersect U.

We now show that western-most points can be determined locally. We call a point p in the closure of a subset C of S^{d-1} a *local optimum* of C if p is the western-most point of the intersection of C with some open subset of S^{d-1} containing p.

Lemma 7 *A point p is a local optimum of a cell C of a Gaussian map G if and only if it is a western-most point of C.*

The proof is in Sect. 6.1. This lemma is crucial for proving the next result, which establishes that a western-most corner of some cell of $\mathcal{G}(P)$ is also a western-most corner of a cell of the Gaussian map of any partial sum $\mathcal{G}(P_S)$ if and only if its western-most point is a node of $\mathcal{G}(P_S)$.

Lemma 8 *Let W be a western-most corner of a cell C in $\mathcal{G}(P)$, with $\mathcal{N}(F; P)$ the western-most point of C. Then, W is a western-most corner of some cell of the same dimension in the Gaussian map of the partial sum $\mathcal{G}(P_S)$ if and only if $I_F \subseteq S$.*

The proof is in Sect. 6.1. This is the most important lemma. It is the ultimate goal of the definitions in this section, which is to have a witness of the existence of a cell, a witness whose presence in the Gaussian maps of partial sums depends on a simple rule.

However, according to the definitions so far, cells intersecting U do not have a western-most corner. In any cell that intersects U, it is possible to turn around U, always going west, much like the way it is possible to turn around a pole on S^2. To deal with this problem, we consider the restriction of the Gaussian map to U. Let us denote as S_U the intersection of S^{d-1} with U. Then, S_U is the sphere $(0, 0) \times S^{d-3}$, and the restriction of a spherical cell complex on S^{d-1} to S_U also defines a spherical cell complex on S_U. In fact, the restriction to S_U of the Gaussian map on S^{d-1} of a d-dimensional polytope is the Gaussian map on S_U of the orthogonal projection of the polytope onto U.

Since S_U is a sphere affinely isomorphic to S^{d-3}, we can define west on S_U as we do for S^{d-1} (see Fig. 2). For any cell of $\mathcal{G}(P)$ that intersects S_U, we define its western-most corner as the western-most corner of its intersection with S_U in the
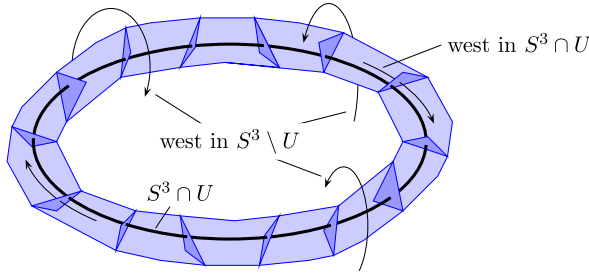
Fig. 2 Representation of a map in S^3 by stereographic projection in Euclidean space. West defined in $S^3 \setminus U$ is turning around the intersection of S^3 with a two-dimensional subspace U, which is affinely isomorphic to S^1. Cells intersecting the subspace have their western-most corner defined by a different direction west, which is the one defined on $S^3 \cap U$

restriction of $\mathcal{G}(P)$ to S_U. If $d > 5$, this again does not define a western-most point for every cell, because west is not defined on the intersection of S_U with a subspace of dimension $d - 4$; so we restrict the Gaussian map to that subspace and start again recursively.

We present now the complete construction. We have chosen a subspace U of dimension $d - 2$, such that its intersection with any two-dimensional plane containing an edge of $\mathcal{G}(P)$ is just the origin. Let us write $U^{d-2} = U$ and denote as G^{d-2} the restriction of the Gaussian map $\mathcal{G}(P)$ to U^{d-2}. G^{d-2} is a spherical cell complex on S^{d-3}. Then, for any i larger than 2, we define from U^i and G^i a subspace U^{i-2}, which is a subspace of U^i, such that the intersection of U^{i-2} with any two-dimensional plane containing an edge of G^i is just the origin. We then define G^{i-2} as the restriction of G^i to U^{i-2}, which is a spherical cell complex on S^{i-3}. This defines a sequence of subspaces $U^{d-2} \supset U^{d-4} \supset \cdots$ and a sequence of spherical cell complexes $G^{d-2} \supset G^{d-4} \supset \cdots$. If d is even, the sequences end with G^2, which is a spherical cell complex on S^1, and U^0 is just the origin and does not intersect S^1. If d is odd, they end with G^3, a spherical cell complex on S^2, and U^1 is a one-dimensional subspace, whose intersection with S^2 defines two antipodal points that we called north and south pole in Sect. 3.

For each spherical cell complex G^i on S^{i-1}, $i \geq 2$, we can define a direction west for every point of S^{i-1} that is not on U^{i-2}, as we have done for $\mathcal{G}(P)$ on S^{d-1}. Then, for any cell C of $\mathcal{G}(P)$, let i be the smallest number such that the intersection of C with U^i is nonempty. We then define the western-most point of C to be the western-most point of $C \cap U^i$, cell of G^i on the sphere S^{i-1}. The western-most corner of C is also the western-most corner of $C \cap U^i$.

Note that if a cell C is of dimension $d - k - 1$, i.e., it is the normal region in $\mathcal{G}(P)$ of a k-dimensional face, it does not intersect U^i for any $i \geq k$, because U^i was chosen so as not to intersect edges and nodes of G^{i+2}, which are restrictions to G^{i+2} of cells of dimension $d - i - 1$ and $d - i - 2$ in $\mathcal{G}(P)$. For instance, if d is odd, only normal region of vertices may intersect with U^1. Since U^1 only intersects S^{d-1} in two antipodal points, there are exactly two cells of dimension $d - 1$ in any Gaussian map that intersect U^1. These are the only two cells that do not have a western-most corner. If d is even, west is defined on every point of G^2, spherical cell complex on S^1, and so every cell of a Gaussian map has a western-most corner.

We have now defined a western-most corner for every cell of Gaussian maps, with the exception, if d is odd, of the two cells that contain a pole. As before, for cells that do not intersect U, the western-most corner of a cell of $\mathcal{G}(P)$ is also a western-most corner of a cell of the same dimension in the Gaussian map of a partial sum $\mathcal{G}(P_S)$ if and only if S contains the support of its western-most point, or rather the support of the cell whose restriction is its western-most point. The cardinality of the support is always less than d.

Now that we have a complete definition of western-most corners, all that remains is to count them. The support of any face of P has cardinality less than d, so all western-most corners of cells of $\mathcal{G}(P)$ can be found in the Gaussian map of partial sums of at most $d-1$ summands. It is not difficult to see that for any $j \geq |I_F|$, there are $\binom{r-|I_F|}{j-|I_F|}$ subsets of $\{1, \ldots, r\}$ of cardinality j that contain I_F. The formula of Theorem 1 is based on the following combinatorial equivalence:

Lemma 9 *For any* $1 \leq s < d \leq r$,

$$\sum_{j=1}^{d-1}(-1)^{d-1-j}\binom{r-1-j}{d-1-j}\binom{r-s}{j-s}=1.$$

The proof is in Sect. 6.2. By this lemma, if we count all the western-most corners in partial sums of j polytopes, multiply by $(-1)^{d-1-j}\binom{r-1-j}{d-1-j}$, and sum over j, we end up counting exactly once each western-most corner, no matter what is the cardinality of the relevant support. Therefore, if $w_k(g)$ is the number of western-most corners of k-dimensional cells in a Gaussian map g, then

$$w_k\big(\mathcal{G}(P)\big)=\sum_{j=1}^{d-1}(-1)^{d-1-j}\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}w_k\big(\mathcal{G}(P_S)\big),\quad k=0,\ldots,d-1,$$

where $\mathcal{C}_j^r$ is the family of subsets of $\{1, \ldots, r\}$ of cardinality j. Since there is one western-most corner of a k-dimensional cell for each $d-1-k$ face of the underlying polytope, this proves Theorem 1 for any d and k. The only exception is that if d is odd, any Gaussian map has two regions of dimension $d-1$ that contain the poles and that have no western-most corner. In that case, $w_{d-1}(\mathcal{G}(\mathcal{P}))=f_0(\mathcal{P})-2$, which gives the special case of the theorem for d odd and $k=0$. Furthermore, each western-most corner of $\mathcal{G}(P)$ can be found at least once in Gaussian maps of partial sums of $d-1$ summands. The last term of the sum in Theorem 1 is thus an upper bound on the whole sum, which proves Corollary 2.   □

5 Maximum Number of Vertices

In this section, we first present a bound on the number of vertices of Minkowski sums using Corollary 2. We then prove Theorem 3 by presenting a construction of Minkowski sums that maximizes the number of vertices of the sum for given number of vertices of the summands.

The trivial bound tells us that if $r < d$, then $f_0(P_1 + \cdots + P_r) \leq \prod_{i=1}^{r} f_0(P_i)$. Consequently, if $r \geq d$, we get by Corollary 2 that

$$f_0(P_1 + \cdots + P_r) \leq \sum_{S \in \mathcal{C}_{d-1}^{r}} \prod_{i \in S} f_0(P_i).$$

This can be seen as enumerating all possible combinations of $d - 1$ vertices chosen each from a different summand. This is necessarily lower than all possible combinations of $d - 1$ vertices from the summands. If the summands have n vertices in total, this upper bound is $\binom{n}{d-1}$, which is in $O(n^{d-1})$. If each summand has n vertices, then we have

$$f_0(P_1 + \cdots + P_r) \leq \sum_{S \in \mathcal{C}_{d-1}^{r}} n^{d-1} = \binom{r}{d-1} n^{d-1}.$$

Proof of Theorem 3 We give a construction that reaches the indicated bound, then prove that no sum can have more vertices. A construction from [2] allows us to choose $d - 1$ polytopes (full-dimensional after perturbation) such that the number of vertices of their sum reaches the trivial bound. We adapt this construction so as to choose r polytopes such that *any* partial sum of $d - 1$ summands reaches the trivial bound.

The previous construction builds $d - 1$ polytopes $P_1, \ldots, P_{d-1}$ whose sum has $\prod_{i=1}^{d-1} f_0(P_i)$ vertices. In that construction, each P_i is the convex hull of $f_0(P_i)$ vertices disposed in an arc of a circle in the space spanned by $\mathbf{e_i}$ and $\mathbf{e_d}$, where $\mathbf{e_1}, \ldots, \mathbf{e_d}$ form an orthonormal basis of the d-dimensional space. The fact that the basis is orthonormal is not actually necessary; the construction also works for any linearly independent vectors $\mathbf{e_1}, \ldots, \mathbf{e_{d-1}}$ orthogonal to $\mathbf{e_d}$. The resulting polytope is then simply the image of the original construction through a nonsingular linear transformation and has the same number of vertices.

Let us therefore choose r vectors $\mathbf{f_1}, \ldots, \mathbf{f_r}$ in the space orthogonal to $\mathbf{e_d}$, such that any subset of $d - 1$ of these r vectors is linearly independent. That is, we choose r vectors in generic position. We then build each polytope P_i, $i = 1, \ldots, r$, similarly to the construction from [2] in the space spanned by $\mathbf{f_i}$ and $\mathbf{e_d}$. If we add a small perturbation on the vertices, we ensure that each polytope is full-dimensional and also that the polytopes are in general orientations. We get that the number of vertices of any partial sum P_S of $d - 1$ of these polytopes is $f_0(P_S) = \prod_{i \in S} f_0(P_i)$. By the trivial bound, this implies that the number of vertices of any partial sum P_S of *less than d* of the polytopes is $f_0(P_S) = \prod_{i \in S} f_0(P_i)$. And so, by Theorem 1, the construction reaches the bound of Theorem 3.

Finally, we prove that for a given number of vertices of the summands, no Minkowski sum of full-dimensional polytopes can have more vertices. We know that the maximum number of vertices is attained by polytopes in general orientations, so we can assume that the polytopes are in general orientations. The number of vertices of their Minkowski sum is given by the formula from Theorem 1.

We then show that this formula never reaches values greater than the values reached when applied to our construction. We use for this the following lemma, which shows that for any full-dimensional polytopes in general orientations, our construction maximizes the sum of two adjacent terms in the alternating sum of the formula.

Lemma 10 *Let $P_1, \ldots, P_r$ be full-dimensional polytopes in general orientations. Then, for any $j > 1$,*

$$\left(\binom{r-1-j}{d-1-j} \sum_{S \in \mathcal{C}^r_j} f_0(P_S) - \binom{r-j}{d-j} \sum_{S \in \mathcal{C}^r_{j-1}} f_0(P_S) \right)$$

$$\leq \left(\binom{r-1-j}{d-1-j} \sum_{S \in \mathcal{C}^r_j} \left(\prod_{i \in S} f_0(P_i) \right) - \binom{r-j}{d-j} \sum_{S \in \mathcal{C}^r_{j-1}} \left(\prod_{i \in S} f_0(P_i) \right) \right).$$

The proof is quite technical. It appears in Sect. 6.3.

The formula from Theorem 1 can be decomposed into such pairs of terms. If d is even, the term for $j = 1$ is not in a pair, but it is constant, since the number of vertices of the summands is fixed. If d is odd, the formula is modified by the value α, but the total value it adds to the formula is also constant. Given that the formula can be decomposed into terms maximized by our construction, our construction maximizes the formula. This concludes the proof of Theorem 3. □

Note, for example, that we cannot extend this result to facets. The maximum number of facets of a Minkowski sum of polytopes is only known for two summands [6], so we cannot write an exact upper bound for facets. However, the complexity of the f-vector is in $O(r^{d-1}n^{d-1})$, as can be deduced by applying Theorem 3.3 of [7] to the upper and lower hull of the dual polytopes of summands.

6 Technical Proofs

6.1 Proof of Lemmas 6, 7, and 8

Recall that by Lemma 5, the subspace U can be chosen so that its intersection with any two-dimensional subspace containing an edge of $\mathcal{G}(P)$ is just the origin. We prove here that U plays the same role as poles in three dimensions and that our definition of west has the property that if we "optimize" in direction west over a spherically convex subset of S^{d-1}, a local optimum of the subset is also a global optimum. We start with a few lemmas:

Lemma 11 *Let p and p' be distinct nonantipodal points of S^{d-1}. Suppose that p and p' are on a same subspace. Then, all points on the great circle of S^{d-1} containing p and p' are on that subspace.*

Proof A great circle of S^{d-1} is the intersection of S^{d-1} with a two-dimensional space. Suppose that a great circle contains p, p', and q, with p and p' on a subspace L, but $q \notin L$. Then, the intersection of the two-dimensional space containing the great circle with L is one-dimensional, and so it is a line going through the origin. Since p and p' are both on that line and in S^{d-1}, they are either the same or antipodal. □

Recall that $\mathbf{e_1}, \ldots, \mathbf{e_d}$ is an orthonormal basis of the d-dimensional space chosen such that $U = span(\{\mathbf{e_3}, \ldots, \mathbf{e_d}\})$. For any θ, let us denote as $L(\theta)$ the subspace orthogonal to $\sin(\theta)\mathbf{e_1} + \cos(\theta)\mathbf{e_2}$, i.e., the set $\{p : \langle p, \sin(\theta)\mathbf{e_1} + \cos(\theta)\mathbf{e_2}\rangle = 0\}$.

Lemma 12 *A point p of S^{d-1} is in $L(\theta) \cap S^{d-1}$ if and only if $\cos(\theta_1(p) - \theta)) = 0$ or $p \in U$.*

Proof Any p in S^{d-1} is written in our parameterization as $\rho s + u$, with $\rho \geq 0$, $s \in S^1$, and $u \in U$. We can write $s = \sin(\theta_1(p))\mathbf{e_1} + \cos(\theta_1(p))\mathbf{e_2}$, and $\rho = 0$ if and only if $p \in U$. Therefore, $\langle p, \sin(\theta)\mathbf{e_1} + \cos(\theta)\mathbf{e_2}\rangle = \rho(\sin(\theta)\sin(\theta_1(p)) + \cos(\theta)\cos(\theta_1(p))) = \rho\cos(\theta_1(p) - \theta)$. So p is in L if and only if $\rho\cos(\theta_1(p) - \theta) = 0$, which is if and only if $\rho = 0$ or $\cos(\theta_1(p) - \theta) = 0$. The result follows. $\qquad\square$

Lemma 13 *Let K be great circle of S^{d-1}. Then, either K is inside U; or K intersects U, $K \setminus U$ has two connected components K_1 and K_2, such that for any two points $p \in K_1$, $p' \in K_2$, $\theta_1(p) + \pi = \theta_1(p')$; or K does not intersect U, and for any two distinct points p, p' in K, $\theta_1(p) \neq \theta_1(p')$, and $\theta_1(p) + \pi = \theta_1(p')$ if and only if p and p' are antipodal.*

Proof Suppose p and p' distinct in K such that $\theta_1(p) = \theta_1(p')$; or suppose p and p' nonantipodal in K such that $\theta_1(p) + \pi = \theta_1(p')$; or suppose p' is in $K \cap U$ and any p in K. In all three cases, by Lemma 12, p and p' are both in the subspace $L(\theta_1(p) + \pi/2)$. By Lemma 11, all points on K are in $L(\theta_1(p) + \pi/2)$. Then, by Lemma 12 again, for any q on the arc of great circle, $q \in U$ or $\cos(\theta_1(q) - (\theta_1(p) + \pi/2)) = 0$. Suppose that $K \cap U$ contains more than two points. Then, some of them are distinct and nonantipodal, and by Lemma 11, K is inside U. Otherwise, for any q and q' antipodal on $K \setminus U$, $\theta_1(q') = \theta_1(q) \pm \pi$. So there must be two antipodal points of K inside U separating q and q', and so $K \setminus U$ has two connected components.

The only remaining case is that for any p and p' distinct in K, $\theta_1(p) \neq \theta_1(p')$; for any p and p' in K such that $\theta_1(p) + \pi = \theta_1(p')$, p and p' must be antipodal; and $K \cap U$ is empty. $\qquad\square$

Note that if a great circle of S^{d-1} does not intersect U, then θ_1 is different in any two points of the great circle. Since the parameterization is smooth on $S^{d-1} \setminus U$, θ_1 augments monotonically and continuously in one direction around the great circle, except in one point when it drops from 2π to 0.

We now recall Lemma 6 and prove it.

Lemma 6 *If a spherically convex subset C of S^{d-1} does not intersect U, it is in a hemisphere defined by $\theta_1 \in [\alpha, \alpha + \pi]$ or $\theta_1 \in [0, \alpha] \cup [\alpha + \pi, 2\pi)]$ for some $\alpha \in [0, 2\pi)$.*

Proof Let $T_1(C)$ be the set of values of θ_1 over C in our parameterization. Since C does not intersect U, $T_1(C)$ is connected. Suppose $T_1(C)$ is $[0; 2\pi)$, then there are two points p, p' in C with $\theta_1(p) + \pi = \theta_1(p')$. Because C is spherically convex, any shortest arc of great circle between p and p' is contained in C. If p and p' are

antipodal, then C is the whole sphere and intersects U. Otherwise, by Lemma 13, the arc of great circle again contains a point in U. This is a contradiction.

Otherwise, suppose without loss of generality that the supremum of $T_1(C)$ is $3\pi/2$. Then, either C is in the hemisphere defined by $\theta_1 \in [\pi/2, 3\pi/2]$, or there is a $\delta > 0$ such that there are two points p, p' in C with $\theta_1(p) + \pi = \theta_1(p') = 3/2 - \delta$. As above, this implies that C contains a point in U, which is a contradiction. $\square$

Let us recall Lemma 7 before proving it.

Lemma 7 *A point p is a local optimum of a cell C of a Gaussian map G if and only if it is a western-most point of C.*

Proof By definition, a western-most point is always a local optimum. Assume that p is a local optimum of C, \ and that some distinct p' is the western-most point of C and therefore also a local optimum. Then, the shortest arc of great circle between p and p' is in C. Let $\alpha = \theta_1(p') - \theta_1(p)$. If $\cos(\alpha) \neq 0$, then by Lemma 13, the great circle defined by p and p' does not intersect U, and θ_1 augments continuously from p to p' except possibly in one point when it jumps from 2π to 0. So there is a q in the intersection of the arc of great circle from p to p' with the open set that proves p is a local optimum. For q close enough, $\theta_1(q) > \theta_1(p)$, and so p is not a local optimum, a contradiction.

Suppose now that $\alpha = \pi$. If p and p' are antipodal, then any great circle containing p and p' is in C, and C is the whole sphere, a contradiction. If p and p' are not antipodal, then by Lemma 13, there is a point in the arc of great circle from p to p' that is in U, and so C intersects U. But this means that C has no western-most point, a contradiction.

Suppose now that $\alpha = 0$. Then, p and p' are incident to a cell, where θ_1 is fixed. But this means that the great circles containing edges of the cell intersect U, which contradicts the way we have chosen U.

Therefore, it is impossible to have a local optimum p and a distinct western-most point p' of a same cell. $\square$

We need to introduce one more lemma before proving Lemma 8. Recall that $\varepsilon > 0$ is smaller than the distance between any two nonincident cells in $\mathcal{G}(P)$.

Lemma 14 *Let F be a facet of P, with its normal region $\mathcal{N}(F; P)$ a node of $\mathcal{G}(P)$. Let p be a point of S^{d-1} at distance less than ε of $\mathcal{N}(F; P)$. For any partial sum P_S with $I_F \subseteq S$, the dimensions of the cells containing p in $\mathcal{G}(P)$ and $\mathcal{G}(P_S)$ are the same.*

Proof In the Gaussian map $\mathcal{G}(P)$, the subset of S^{d-1} at a distance less than ε of $\mathcal{N}(F; P)$ intersects only the normal regions of subfaces of F. Therefore, for any point p in that subset, $\mathcal{S}(P; p)$ is a subface G of F. Recall that for any subface G of a facet F, $I_G \subseteq I_F$. So for any partial sum P_S such that $I_F \subseteq S$, $I_G \subseteq S$, which means that not only P_S has a facet with the same normal region as F, but $\mathcal{S}(P_S; p)$ is a subface of that facet with the same dimension as G, and p is in a cell of the same dimension in $\mathcal{G}(P_S)$ as in $\mathcal{G}(P)$. $\square$

We finally have the tools to prove Lemma 8.

Lemma 8 *Let W be a western-most corner of a cell C in $\mathcal{G}(P)$, with $\mathcal{N}(F; P)$ the western-most point of C. Then, W is a western-most corner of some cell of the same dimension in the Gaussian map of the partial sum $\mathcal{G}(P_S)$ if and only if $I_F \subseteq S$.*

Proof First, if $I_F \nsubseteq S$, then $\mathcal{N}(F; P)$ is not a node of $\mathcal{G}(P_S)$, and so W cannot be a western-most corner. Suppose $I_F \subseteq S$; then $\mathcal{N}(F; P)$ is a node of $\mathcal{G}(P_S)$. Furthermore, by Lemma 14, the points in W are in a cell of the same dimension in $\mathcal{G}(P_S)$ as in $\mathcal{G}(P)$, and the points in the closure of W are in a cell of the same dimension in $\mathcal{G}(P_S)$ as in $\mathcal{G}(P)$. As a consequence, since $\mathcal{N}(F; P)$ is the western-most point of W in $\mathcal{G}(P)$, it is also the western-most point of W in $\mathcal{G}(P_S)$. But W is the intersection, of the cell it is in, with an open subset, and so $\mathcal{N}(F; P)$ is a local optimum of the cell that contains W in $\mathcal{G}(P_S)$. By Lemma 7, it is also the western-most point of the cell that contains W in $\mathcal{G}(P_S)$, and so W is the western-most corner of that cell in $\mathcal{G}(P_S)$. □

6.2 Proof of Lemma 9

We prove here the combinatorial equivalence used for formulating Theorem 1. The relation can also be derived from a protean family of equivalences of the form

$$\sum_{j=0}^{c} (-1)^{c-j} \binom{a+j}{b+c} \binom{c}{j} = \binom{a}{b}, \quad b < a, c \geq 0.$$

See also [13, p. 285] on this subject. Let us recall Lemma 9 before proving it.

Lemma 9 *For any $1 \leq s < d \leq r$,*

$$\sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{r-1-j}{d-1-j} \binom{r-s}{j-s} = 1.$$

Proof We prove the lemma by induction over r. We know that for any $d > 1$, $\sum_{j=0}^{d} (-1)^j \binom{d}{j} = 0$. We can also write $\sum_{j=1}^{d} (-1)^j \binom{d-s}{j-s} = 0$ for any $1 \leq s < d$, and so we have $\sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{d-s}{j-s} = 1$. We can also write

$$\sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{d-1-j}{d-1-j} \binom{d-s}{j-s} = 1.$$

This proves the relation for $r = d$. Assume that the relation is proved for r. Then,

$$\sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{r-1-j}{d-1-j} \binom{r-s}{j-s} = 1,$$

$$= \sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{r-j}{d-1-j} \binom{r-s}{j-s} - \sum_{j=1}^{d-1} (-1)^{d-1-j} \binom{r-1-j}{d-2-j} \binom{r-s}{j-s}.$$

Replacing j in the second sum with $j' - 1$, we get

$$\sum_{j=1}^{d-1}(-1)^{d-1-j}\binom{r-j}{d-1-j}\binom{r-s}{j-s}$$

$$-\sum_{j'=2}^{d}(-1)^{d-j'}\binom{r-j'}{d-1-j'}\binom{r-s}{j'-1-s} = 1.$$

In the second sum, the term $j' = d$ gives zero, so we can remove it and add one for $j' = 1$, which also gives zero:

$$\sum_{j=1}^{d-1}(-1)^{d-1-j}\binom{r-j}{d-1-j}\binom{r-s}{j-s}$$

$$-\sum_{j'=1}^{d-1}(-1)^{d-j'}\binom{r-j'}{d-1-j'}\binom{r-s}{j'-1-s} = 1.$$

Grouping the sums, we get

$$\sum_{j=1}^{d-1}(-1)^{d-1-j}\binom{(r+1)-1-j}{d-1-j}\binom{(r+1)-s}{j-s} = 1,$$

and so the relation holds for $r + 1$, which proves Lemma 9 by induction. $\qquad\square$

6.3 Proof of Theorem 3

We prove here Theorem 3, by providing a proof to Lemma 10.

The formula from Theorem 1 that gives the number of vertices of full-dimensional polytopes in general orientations can be decomposed into pairs of adjacent terms, plus a constant term. If the sum of each pair is maximal when using our construction, the sum of all pairs is maximal as well, and so the expression obtains its maximal value when using our construction. This is shown using Lemma 10 that we recall:

Lemma 10 *Let $P_1, \ldots, P_r$ be full-dimensional polytopes in general orientations. Then, for any $j > 1$,*

$$\left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}f_0(P_S) - \binom{r-j}{d-j}\sum_{S\in\mathcal{C}_{j-1}^r}f_0(P_S)\right)$$

$$\leq \left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}\left(\prod_{i\in S}f_0(P_i)\right) - \binom{r-j}{d-j}\sum_{S\in\mathcal{C}_{j-1}^r}\left(\prod_{i\in S}f_0(P_i)\right)\right).$$

Proof First, we represent the internal sum $\sum_{S\in\mathcal{C}_{j-1}^r}f_0(P_S)$ as a sum over $S \in \mathcal{C}_j^r$:

$$\sum_{S\in\mathcal{C}_{j-1}^r}f_0(P_S) = \frac{1}{r-j+1}\sum_{S\in\mathcal{C}_j^r}\sum_{i\in S}f_0(P_{S\setminus i}).$$

That is, subsets in $\mathcal{C}_{j-1}^r$ are obtained $r - j + 1$ times by enumerating subsets in $\mathcal{C}_j^r$ and removing each element in turn. By the trivial bound, $f_0(P_{S\setminus i}) \geq \frac{f_0(P_S)}{f_0(P_i)}$. Since $\binom{r-j}{d-j} = \frac{r-j}{d-j}\binom{r-1-j}{d-1-j}$, we have

$$\left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}f_0(P_S) - \binom{r-j}{d-j}\sum_{S\in\mathcal{C}_{j-1}^r}f_0(P_S)\right)$$

$$= \left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}\left(f_0(P_S) - \frac{r-j}{d-j}\frac{1}{r-j+1}\sum_{i\in S}f_0(P_{S\setminus i})\right)\right)$$

$$\leq \left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}\left(f_0(P_S)\left(1 - \frac{r-j}{d-j}\frac{1}{r-j+1}\sum_{i\in S}\frac{1}{f_0(P_i)}\right)\right)\right).$$

All summands P_i are full-dimensional, and so $f_0(P_i) \geq d + 1$. Since S has j elements, we have

$$\left(1 - \frac{r-j}{d-j}\frac{1}{r-j+1}\sum_{i\in S}\frac{1}{f_0(P_i)}\right) \geq \left(1 - \frac{1}{d-j}\frac{j}{d+1}\right) \geq 0.$$

So each term is positive. Therefore, the trivial bound $f_0(P_S) \leq \prod_{i\in S}f_0(P_i)$ implies that

$$\left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}f_0(P_S) - \binom{r-j}{d-j}\sum_{S\in\mathcal{C}_{j-1}^r}f_0(P_S)\right)$$

$$\leq \left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}\left(\prod_{i\in S}f_0(P_i)\left(1 - \frac{r-j}{d-j}\frac{1}{r-j+1}\sum_{i\in S}\frac{1}{f_0(P_i)}\right)\right)\right).$$

Now, since we have

$$\sum_{S\in\mathcal{C}_j^r}\left(\prod_{i\in S}f_0(P_i)\right)\left(\sum_{i\in S}\frac{1}{f_0(P_i)}\right) = \sum_{S\in\mathcal{C}_j^r}\sum_{i\in S}\prod_{i'\in S\setminus i}f_0(P_{i'})$$

$$= (r-j+1)\sum_{S\in\mathcal{C}_{j-1}^r}\prod_{i\in S}f_0(P_i),$$

we finally get that

$$\left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}f_0(P_S) - \binom{r-j}{d-j}\sum_{S\in\mathcal{C}_{j-1}^r}f_0(P_S)\right)$$

$$\leq \left(\binom{r-1-j}{d-1-j}\sum_{S\in\mathcal{C}_j^r}\left(\prod_{i\in S}f_0(P_i)\right) - \binom{r-j}{d-j}\sum_{S\in\mathcal{C}_{j-1}^r}\left(\prod_{i\in S}f_0(P_i)\right)\right). \qquad \square$$

The above proves that for any choice of polytopes $P_1, \ldots, P_r$, the sum of each pair of terms in the formula of Theorem 1 is not greater than the term we obtain with

polytopes of our construction, where we have $f_0(P_S) = \prod_{i \in S} f_0(P_i)$. Therefore, for fixed number of vertices of the summands, our construction gives a Minkowski sum with the maximum number of vertices.

7 Summary

We extend the intuitive concept of west from three dimensions to higher dimensions. The important property of this concept is that in a spherically convex region, any local western-most point is also a global western-most point. This enables us to prove a relation on the number of faces of sums of many polytopes in general orientations.

We use this relation to provide an upper bound on the maximum number of vertices of Minkowski sums for given number of vertices of summands, and we show that this bound can be reached.

References

1. Fogel, E., Halperin, D., Weibel, C.: On the exact maximum complexity of Minkowski sums of polytopes. Discrete Comput. Geom. **42**(4), 654–669 (2009)
2. Fukuda, K., Weibel, C.: On f-vectors of Minkowski additions of convex polytopes. Discrete Comput. Geom. **37**, 503–516 (2007)
3. Fukuda, K., Weibel, C.: A linear equation for Minkowski sums of polytopes relatively in general position. Eur. J. Comb. **31**(2), 565–573 (2010). Combinatorics and geometry
4. Gritzmann, P., Sturmfels, B.: Minkowski addition of polytopes: computational complexity and applications to Gröbner bases. SIAM J. Discrete Math. **6**(2), 246–269 (1993)
5. Halperin, D., Kavraki, L.E., Latombe, J.-C.: Robotics. In: Goodman, J.E., O'Rourke, J. (eds.) Handbook of Discrete and Computational Geometry, Chap. 48, pp. 1065–1093. CRC Press, Boca Raton (2004)
6. Karavelas, M., Tzanaki, E.: The maximum number of faces of the Minkowski sum of two convex polytopes. In: ACM-SIAM Symposium on Discrete Algorithms (2012)
7. Koltun, V., Sharir, M.: On overlays and minimization diagrams. Discrete Comput. Geom. **41**(3), 385–397 (2009)
8. Lozano-Pérez, T., Wesley, M.A.: An algorithm for planning collision-free paths among polyhedral obstacles. Commun. ACM **22**(10), 560–570 (1979)
9. Pachter, L., Sturmfels, B. (eds.): Algebraic Statistics for Computational Biology. Cambridge University Press, New York (2005)
10. Sanyal, R.: Topological obstructions for vertex numbers of Minkowski sums. J. Comb. Theory, Ser. A **116**(1), 168–179 (2009)
11. Weibel, C.: Minkowski sums of polytopes: combinatorics and computation. PhD thesis, EPFL, Lausanne (2007)
12. Zhang, H.: Partially observable Markov decision processes: a geometric technique and analysis. Oper. Res., doi:10.1287/opre.1090.0697 (2009)
13. Ziegler, G.M.: Lectures on Polytopes. Graduate Texts in Mathematics, vol. 152. Springer, New York (1995)