Embedding Stacked Polytopes on a Polynomial-Size Grid

Erik D. Demaine* André Schulz†

November 9, 2018

Abstract

A stacking operation adds a d-simplex on top of a facet of a simplicial d-polytope while maintaining the convexity of the polytope. A stacked d-polytope is a polytope that is obtained from a d-simplex and a series of stacking operations. We show that for a fixed d every stacked d-polytope with n vertices can be realized with nonnegative integer coordinates. The coordinates are bounded by $O(n^{2\log_2(2d)})$, except for one axis, where the coordinates are bounded by $O(n^{3\log_2(2d)})$. The described realization can be computed with an easy algorithm.

The realization of the polytopes is obtained with a lifting technique which produces an embedding on a large grid. We establish a rounding scheme that places the vertices on a sparser grid, while maintaining the convexity of the embedding.

1 Introduction

Steinitz's Theorem [Ste22, Zie95] states that the graphs of 3-polytopes[1] are exactly the planar 3-connected graphs. In particular, every planar 3-connected graph can be realized as a 3-polytope. The original proof is constructive, transforming the graph by a sequence of local operations down to a tetrahedron. Unfortunately, the resulting polytope construction requires exponentially many bits of accuracy for each vertex coordinate. Stated another way, this construction can place the n vertices on an integer grid, but that grid may have dimensions doubly exponential in n [OS94]. The situation in higher dimensions is even worse. Already in dimension 4, there are polytopes that cannot be realized with rational coordinates, and a 4-polytope that can be realized on the grid might require coordinates that are doubly exponential in the number of its vertices [Zie95]. Moreover, it is NP-hard to decide whether a lattice is a face lattice of a 4-polytope [RG96, RGZ95].

How large an integer grid do we need to embed a given planar 3-connected graph with n vertices as a polytope? This question goes back at least eighteen years as Problem 4.16 in Günter M. Ziegler's book [Zie95]; he wrote that "it is quite possible that there is a quadratic upper bound" on the length of the maximum dimension. The best bound so far is exponential in n, namely $O(2^{7.21n})$ [BS10, MRS11]; see below for the long history. The central question is whether a polynomial grid suffices, that is, whether Steinitz's Theorem can be made efficient. For comparison, a planar graph can be embedded in the plane with strictly convex faces using a polynomial-size grid [BR06]. In this paper, we give the first nontrivial subexponential upper bound for a large class of 3-polytopes. Moreover, our construction generalizes to higher dimensions and we show that a nontrivial class of d-polytopes can be realized with integers coordinates, which are bounded by a polynomial in n.

A *d-dimensional stacked polytope* is a polytope that is constructed by a sequence of "stacking operations" applied to a d-simplex. A *stacking operation* glues a d-simplex Δ atop a simplicial facet f of polytope, by identifying f with a face of Δ, while maintaining the convexity of the polytope. Thus a stacking operation removes one facet f and adds d new facets having a new common vertex. We call this new vertex *stacked on f*.

*Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA. edemaine@mit.edu
†LG Theoretische Informatik, FernUniversität in Hagen, andre.schulz@fernuni-hagen.de. Supported by the German Research Foundation (DFG) under grant SCHU 2458/2-1.

[1] In our terminology a *polytope* is always understood as a *convex* polytope.

Our results. We present an algorithm that realizes a stacked polytope on a grid whose dimensions are polynomial in n. In our presentation log denotes the binary logarithm. Our main result is the following:

Theorem 1 *Every d-dimensional stacked polytope can be realized on an integer grid, such that all coordinates have size at most* $10d^2R^2$*, except for one axis, where the coordinates have size at most* $6R^3$*, for* $R = dn^{\log(2d)}$*.*

As a corollary of Theorem 1 we obtain that every stacked 3-polytope can be embedded on the grid of dimensions $270n^{5.17} \times 270n^{5.17} \times 18n^{7.76}$.

Related work. Several algorithms have been developed to realize a given graph as a 3-polytope. Most of these algorithms are based on the following two-stage approach. The first stage computes a plane (*flat*) embedding. To extend the plane drawing to a 3-polytope, the plane drawing must fulfill a criterion which can be phrased as an "equilibrium stress condition". Roughly speaking, replacing every edge of the graph with a spring, the resulting system of springs must be in a stable state for the plane embedding. Plane drawings that fulfill this criterion for the interior vertices can be computed as barycentric embeddings, i.e., by Tutte's method [Tut60, Tut63]. The main difficulty is to guarantee the equilibrium condition for the boundary vertices as well, because in general this goal is achievable only for certain locations of the outer face. The second stage computes a 3-polytope by assigning every vertex a height expressed in terms of the spring constants of the system of springs.

The two-stage approach finds application in a series of algorithms [CGT96, EG95, HK92, OS94, MRS11, RG96, Sch11]. The first result that improves the induced grid embedding of Steinitz's construction is due to Onn and Sturmfels [OS94]; they achieved a grid size of $O(n^{160n^3})$. Richter-Gebert's algorithm [RG96] uses a grid of size $O(2^{18n^2})$ for general 3-polytopes, and a grid of size $O(2^{5.43n})$ if the graph of the polytope contains at least one triangle. These bounds were improved by Ribó Mor [Rib06] and later on by Ribó Mor, Rote, and Schulz [MRS11]. The last paper expresses an upper bound for the grid size in terms of the number of spanning trees of the graph. Using the recent bounds of Buchin and Schulz [BS10] on the number of spanning trees, this approach gives an upper bound on the grid size of $O(2^{7.21n})$ for general 3-polytopes and $O(2^{4.83n})$ for 3-polytopes with at least one triangular face. These bounds are the best known to date for the general case. Very recently, Pak and Wilson proved that every simplicial[2] 3-polytope can be embedded on a grid of size $4n^3 \times 8n^5 \times (500n^8)^n$ [PW13].

Zickfeld showed in his PhD thesis [Zic07] that it is possible to embed very special cases of stacked 3-polytopes on a grid polynomial in n. First, if each stacking operation takes place on one of the three faces that were just created by the previous stacking operation (what might be called a *serpentine stacked polytope*), then there is an embedding on the $n \times n \times 3n^4$ grid. Second, if we perform the stacking in rounds, and in every round we stack on every face simultaneously (what might be called the *balanced* stacked polytope), then there is an embedding on a $\frac{4}{3}n \times \frac{4}{3}n \times O(n^{2.47})$ grid. Zickfeld's embedding algorithm for balanced stacked 3-polytopes constructs a barycentric embedding. Because of the special structure of the underlying graph the plane embedding remarkably fits on a small grid.

Every stacked polytope can be extended to a balanced stacked polytope at the expense of adding an exponential number of vertices. By doing so, Zickfeld's grid embedding for the balanced case induces a $O(2^{3.91n})$ grid embedding for general stacked 3-polytopes.

Little is known about the lower bound of the grid size for embeddings of 3-polytopes. An integral convex embedding of an n-gon in the plane needs an area of $\Omega(n^3)$ [AŽ95, And61, Thi91]. Therefore, realizing a 3-polytope with an $(n-1)$-gonal face requires at least one dimension of size $\Omega(n^{3/2})$. For simplicial polytopes (and hence for stacked polytopes), this lower-bound argument does not apply. However, there are planar 3-trees that require $\Omega(n^2)$ area for plane straight-line drawings [MNRA11]. We show in Sect. 4 how to get a lower bound of $\Omega(n^3)$ volume for stacked 3-polytopes based on this 2d example.

[2] A polytope is *simplicial* if all its faces are simplices.

Alternative approaches for realizing general 3-polytopes come from the original proof of Steinitz's theorem, as well as the Koebe–Andreev–Thurston circle-packing theorem, which induces a particular polytope realization called the *canonical polytope* [Zie95]. Das and Goodrich [DG97] essentially perform many inverted edge contractions on many independent vertices in one step, resulting in a singly exponential bound on the grid size. The proof of the Koebe–Andreev–Thurston circle-packing theorem relies on nonlinear methods and makes the features of the 3-dimensional embedding obtained from a circle packing intractable; see [Sch91] for an overview. Lovász studied a method for realizing 3-polytopes using a vector of the nullspace of a *Colin de Verdière matrix* of rank 3 [Lov00]. It is easy to construct these matrices for stacked polytopes; however, without additional requirements, the computed grid embedding might again need an exponential-size grid.

Our contribution. At a high level, we follow the popular two-stage approach: we compute a *flat embedding* in $\mathbf{R}^{d-1}$ and then lift it to a d-polytope. Although this two-stage approach is well known for realizing 3-polytopes, it cannot be easily extended to higher dimensions. There exists a generalization of equilibrium stresses for polyhedral complexes in higher dimensions by Rybnikov [Ryb99]. However, it is not straightforward to operate with the formalism used by Rybnikov for our purposes. To overcome the difficulties of handling the more complex behavior of the higher-dimensional polytopes and to keep our presentation self-contained, we develop our own specialized methods to study liftings of *stacked* polytopes. Notice that our methods coincides with the approach of Rybnikov, however, we omit the proof, since it is not required for this presentation.

Instead of specifying the stress and then computing the barycentric embedding we construct the "stress" and the barycentric embedding in parallel. To specify the stress we define the heights of the lifting (actually, the vertical movement of the vertices as induced by the stacking operation). In our presentation the concept of stress is therefore reduced to a certificate for the convexity of the lifting, but it is not defining the flat embedding. On the other hand we still use barycentric coordinates to determine the flat embedding. A crucial step in our algorithm is the construction of a balanced set of barycentric coordinates, which corresponds to face volumes in the flat embedding. Initially, all faces have the same volume, but to prevent large heights in the lifting, we increase the volumes of the small faces. To see which faces must be blown up, we make use of a decomposition technique from data structural analysis called *heavy path decomposition* [Tar83]. Based on this decomposition, we subdivide the stacked polytope into a hierarchy of (serpentine) stacked polytopes, which we use to define the barycentric coordinates. At this stage the lifting of the flat embedding would result in a grid embedding with exponential coordinates. But since we have balanced the volume assignments, we can allow a small perturbation of the embedding, while maintaining its convexity. Analyzing the size of the feasible perturbations shows that we can round to points on a polynomially sized grid.

A preliminary version of this work was presented at the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA) in San Francisco [DS11]. In the preliminary version we were focused on the more prominent 3-dimensional case and did not present any bounds for higher-dimensional polytopes. For the sake of a unified presentation we changed the construction of the lifting slightly. In the preliminary version we defined the constructed stress as a linear combination of stresses defined on certain K_4s. In this paper we specify the "movement" for every stacked vertex (its vertical shift) directly. This can be considered as the dual definition of the lifting. Another difference is the more careful analysis of the size of the z-coordinates in the final embedding. In contrast to the preliminary version, where we presented bound of $224{,}000n^{18}$ for the z-coordinates, we present a different method for bounding the height of the lifting, which yields an upper bound of $18n^{7.76}$ instead. We remark also that the conference version contained a flaw in the area assignment, which is now fixed by a slightly modified construction (Lemma 7). In a paper that followed the preliminary version Igamberdiev and Schulz introduced a duality transformation for 3-polytopes that allows to control the grid size of the dual polytope [IS16]. By this our results for stacked polytopes can be transferred to their dual polytopes (truncated 3-polytopes), which shows that also this class can be realized on a polynomial-sized grid.

2 Specifying the geometry of stacked polytopes

In this section we develop the necessary tools for defining an embedding of a stacked d-polytope with the two-stage approach. We present the framework for our algorithm by defining its basic construction and by identifying certificates to guarantee its correctness. The actual algorithm is then presented in Section 3.

2.1 Matrices, determinants, simplices

We start our presentation with introducing some notation. Assume that $S = (\mathbf{s}_1, \ldots, \mathbf{s}_k)$ is a sequence of $(k-1)$-dimensional vectors. Then we denote by $(\mathbf{s}_1, \ldots, \mathbf{s}_k)$ the matrix whose row vectors (in order) are $\mathbf{s}_1, \ldots, \mathbf{s}_k$. We define as

$$[S] = \det \begin{pmatrix} \mathbf{s}_1 & \mathbf{s}_2 & \ldots & \mathbf{s}_k \\ 1 & 1 & \ldots & 1 \end{pmatrix}.$$

Notice that $[S]$ equals $(k-1)!$ times the signed volume of the simplex spanned by $\mathbf{s}_1, \ldots, \mathbf{s}_k$. When working with sequences we use the binary $\circ$ operator to concatenate two sequences. If there is no danger of confusion we identify an element with the singleton set (or sequence) that contains this element. If f is a function on $\mathcal{U}$ we extend f in the natural way to act on sequences on $\mathcal{U}$, i.e., we set $f((a_1, \ldots, a_z)) := (f(a_1), \ldots, f(a_z))$.

Throughout the paper we denote for a sequence X of affinely independent points the simplex spanned by X by Δ_X. If this simplex is the realization of some face of an embedding of a polytope we denote this face by f_X. For a d-polytope we call a $(d-1)$-face a *facet*, and a $(d-2)$-face a *ridge*.

By convention, we understand as $\mathbf{R}^k$ the space spanned by the first k standard basis vectors. In this sense we consider the subspace $\mathbf{R}^{k-1}$ as space embedded inside $\mathbf{R}^k$. Furthermore, when we speak about hyperplanes in the following, we consider only those hyperplanes in $\mathbf{R}^k$ that are not orthogonal to $\mathbf{R}^{k-1}$.

2.2 Flat embedding

An important intermediate step in our embedding algorithm is the construction of a flat embedding, which is a simplicial complex constructed by repeated weighted barycentric subdivisions of a $(d-1)$-simplex Δ_B. The combinatorial structure of the flat embedding represents the face lattice of the stacked polytope. Let Δ_S be a $(d-1)$-simplex and let $\Delta_1, \Delta_2, \ldots, \Delta_d$ be the $(d-2)$-simplices that define its boundary. A *barycentric subdivision* adds a new point $\mathbf{p}$ in the interior of Δ_S and splits Δ_S into d $(d-1)$-simplices. For every simplex Δ_i we obtain a new simplex Δ'_i spanned by $\Delta_i \cup \{\mathbf{p}\}$. If the volume of all Δ'_is are given such that they sum up to the volume of Δ_S, then there is one unique vertex $\mathbf{p}$, such that the subdivision respects this volume assignment. We call the volumes the Δ'_is the *barycentric coordinates* of $\mathbf{p}$ with respect to Δ_S.

We apply repeated barycentric subdivisions to create (a flat) embedding of a stacked polytope. In particular, the subdivisions carry out the stacking operations geometrically (projected in the $z = 0$ hyperplane). We start the subdivision process with two copies of a $(d-1)$-simplex in the $z = 0$ hyperplane glued along the boundary. One of these "initial" simplices will be the *base face* f_B, which remains unaltered during further modifications. The other face will be repeatedly subdivided such that the combinatorial structure of the stacked polytope is obtained. The constructed realization is still flat, which means that it lies in the $z = 0$ hyperplane.

2.3 Liftings

Our algorithm for realizing stacked polytopes is based on the *lifting technique*. The assignment of an additional coordinate to every vertex in the flat embedding is called a *lifting*. For the assignment of a new coordinate z to the point $\mathbf{p} \in \mathbf{R}^{d-1}$ we use the shorthand notation $(\mathbf{p}, z)$. The lifting is complemented by the projection function $\pi \colon \mathbf{R}^d \to \mathbf{R}^{d-1}$ that simply removes the d-th coordinate. As shortcut notation we write $[\![S]\!]$ for the expression $[\pi(S)]$. In the following we refer to the d-th coordinate of a point in $\mathbf{R}^d$ as its z-coordinate.

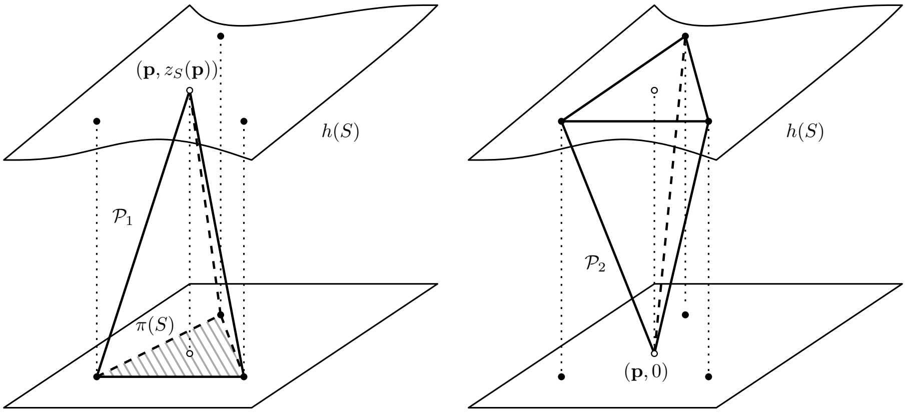

Figure 1: The simplices $\mathcal{P}_1$ and $\mathcal{P}_2$ as defined in the proof of Lemma 1. Both pyramids have the same (absolute) volume.

A hyperplane $h \subset \mathbf{R}^d$ is characterized by a function $z_h \colon \mathbf{R}^{d-1} \to \mathbf{R}$ that assigns to every point $\mathbf{p} \in \mathbf{R}^{d-1}$ a new coordinate, such that $(\mathbf{p}, z_h(\mathbf{p}))$ lies on h. We denote by $h(S)$ the hyperplane *spanned* by S and use as shortcut notation z_S for $z_{h(S)}$. The following lemma gives us a convenient expression for the function z_S.

Lemma 1 *Let S be a set of d affinely independent points from $\mathbf{R}^d$. For every point $\mathbf{p} \in \mathbf{R}^{d-1}$ we have*

$$z_S(\mathbf{p}) = \frac{[S \circ (\mathbf{p}, 0)]}{[\![S]\!]}.$$

Proof. We give a geometric proof of the lemma. Consider the d-simplex $\mathcal{P}_1$ spanned by $\pi(S)$ and $(\mathbf{p}, z_S(\mathbf{p}))$, and the d-simplex $\mathcal{P}_2$ spanned by S and $(\mathbf{p}, 0)$ as depicted in Fig. 1. Simplex $\mathcal{P}_1$ is the affine image of $\mathcal{P}_2$, under the mapping $(\mathbf{q}, z) \mapsto (\mathbf{q}, z_S(\mathbf{q}) - z)$. Since this mapping is a linear shear, it preserves the volumes of the simplices. The simplex $\mathcal{P}_1$ can be understood as a pyramid with base $[\![S]\!]$ and height $z_S(\mathbf{p})$ and therefore its absolute volume *vol* equals

$$vol(\mathcal{P}_1) = \left| \frac{1}{d(d-1)!} [\![S]\!] z_S(\mathbf{p}) \right|.$$

On the other hand we can express the absolute volume of $\mathcal{P}_2$ using the determinant formula and obtain

$$vol(\mathcal{P}_2) = \left| \frac{1}{d!} [S \circ (\mathbf{p}, 0)] \right|.$$

Setting $vol(\mathcal{P}_1) = vol(\mathcal{P}_2)$ proves $z_S(\mathbf{p}) = |[S \circ (\mathbf{p}, 0)]/[\![S]\!]|$. It remains to check if the sign of the expression in the lemma is correct. By the geometric argument the correctness does not depend on the actual configuration given by S and $\mathbf{p}$. To see this note that for a fixed $\mathbf{p}$ a sign change in $[S \circ (\mathbf{p}, 0)]$ occurs if and only if it occurs in $[\![S]\!]$. Moreover, if we keep S fixed, a sign change of $[S \circ (\mathbf{p}, 0)]$ can only happen if $(\mathbf{p}, 0)$ changes its location relative to $h(S)$, which on the other hand also changes the sign for $z_S(\mathbf{p})$. Hence, it suffices to look at a concrete example. Let S be given by the standard basis of $\mathbf{R}^d$, and let $\mathbf{p}$ be the origin in $\mathbf{R}^{d-1}$. In this case, clearly, $[S \circ (\mathbf{p}, 0)] = +1$ and $[\![S]\!] = +1$. Since $z_S(\mathbf{p}) = 1$ the sign is correct. □

2.4 Creases and stresses

Let f, g be two hyperplanes in $\mathbf{R}^d$. Furthermore, let S, resp. T, be a sequence of d affinely independent points of f, resp. g, such that deleting the last element in S and T gives the same subsequence X. It follows that f and g intersect in a flat of dimension $d-2$ that contains the simplex Δ_X. Furthermore, we denote by $\mathbf{s}_d$ the last point in the sequence S, that is the point not in T, and similarly we denote by $\mathbf{t}_d$ the point in T that is not in S. We also set $\mathbf{r} = \pi(\mathbf{s}_d)$. The situation is depicted in Fig. 2. We define as *creasing* on X

$$c(f, g, X) := \frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]}. \tag{1}$$

Roughly speaking, the creasing is a measure for the intersection angle of f and g with respect to the volume of Δ_X.

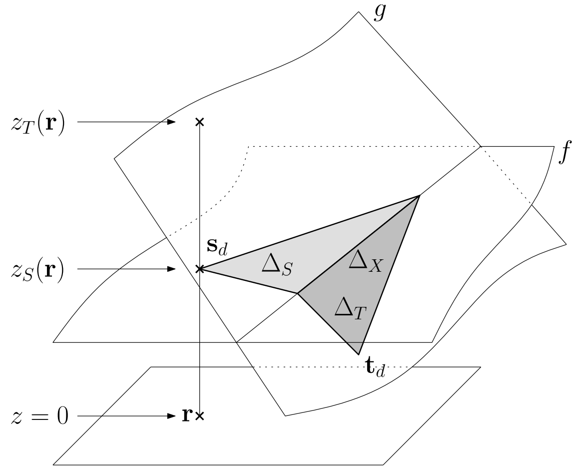

Figure 2: The creasing on X is defined in terms of the objects depicted in the figure (for $d = 3$).

The following lemma shows that the function c is well-defined.

Lemma 2 *Let $c(f, g, X)$ be defined as above, i.e., S spans f, T spans g, such that $S = X \circ \mathbf{s}_d$, and $T = X \circ \mathbf{t}_d$.*

(a) It holds that $c(f, g, X) = -c(g, f, X)$.

(b) The value of $c(f, g, X)$ is independent of the choice of $\mathbf{s}_d$ and $\mathbf{t}_d$.

Proof. Let T be the sequence $(\mathbf{t}_1, \ldots, \mathbf{t}_d)$ and let z_i be the z-coordinate of $\mathbf{t}_i$. We use Lemma 1 to rephrase $c(f, g, X)$ by substituting $z_T(\mathbf{r})$ and obtain

$$c(f, g, X) = \frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]} = \frac{[T \circ (\mathbf{r}, 0)] - z_S(\mathbf{r})[\![T]\!]}{[\![S]\!][\![T]\!]}. \tag{2}$$

To further evaluate the last expression we notice that

$$[T \circ (\mathbf{r}, 0)] - [T \circ (\mathbf{r}, z_S(\mathbf{r}))] = z_S(\mathbf{r})[\![T]\!].$$

Since $[T \circ (\mathbf{r}, z_S(\mathbf{r}))] = [T \circ \mathbf{s}_d]$, we obtain that $[T \circ (\mathbf{r}, 0)] - z_S(\mathbf{r})[\![T]\!] = [T \circ \mathbf{s}_d]$. Plugging the last equation into (2) gives

$$c(f, g, X) = \frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]} = \frac{[T \circ (\mathbf{r}, 0)] - z_S(\mathbf{r})[\![T]\!]}{[\![S]\!][\![T]\!]} = \frac{[T \circ \mathbf{s}_d]}{[\![T]\!][\![S]\!]}. \tag{3}$$

Notice that $[T \circ \mathbf{s}_d] = -[S \circ \mathbf{t}_d]$, since both underlying matrices differ only by one column swap (the sequence of the first $d-1$ members of S and T coincide). As a consequence, exchanging S and T in the right hand side of (3) changes only the sign, which proves statement (a).

To prove statement (b) we argue as follows. Assume we have replaced $\mathbf{t}_d$ in T (that is the only point in T not in X) by some other point on $h(T)$. The replacement of $\mathbf{t}_d$ does not change $z_T(\mathbf{r})$ in Equation (1), and all other parts of the equation only depend on S. Therefore, the expression given in (1) does not depend on $\mathbf{t}_d$. It remains to show, that also changing the last point in S, does not change the creasing of X. However, this follows easily since by statement (a) we can interchange the roles of S and T (this results in a sign change), then we apply the above argument, and then we change S and T back (which cancels the sign change). $\square$

Based on the concept of creasing we are now ready to define the crucial concept of stress. Let us first introduce the notion of right and left facet. We assume that we have given a flat embedding at this point. Let X be the set of vertices of some ridge f_X of P. We fix an arbitrary order for X, and do so also for the vertex sets of the other ridges. The ridge f_X separates two facets of P, say f_S and f_T. As usual, $S = X \cup \mathbf{s}_d$ is the vertex set of f_S. If $[\![X \circ \mathbf{s}_d]\!] > 0$ then f_S is called *left of* f_X, if $[\![X \circ \mathbf{s}_d]\!] < 0$ then f_S is called *right of* f_X. We have one exception from this rule: for the special (base) face f_B the notion of left and right is interchanged. Note that the definitions left/right are independent of the choice of the z-coordinates, and only depend on the flat embedding. Note also, that every ridge has a right and a left incident facet.

Definition 1 (stress) *Let P be stacked d-polytope obtained by lifting a flat embedding and let X be a sequence of vertices which defines a ridge. We define as* stress on X

$$\omega_X := c(h_l, h_r, X), \tag{4}$$

where h_r contains the facet right of X and h_l contains the facet left of X.

We can immediately deduce the following property for the stress along X.

Lemma 3 *The stress on X is not affected by the order of the elements in X.*

Proof. When reordering the elements of X, only the sign of a creasing along X might change. When it flips then our notion of left and right face will also be exchanged, which means that the left face becomes the right face with respect to X and vice versa. Let $\bar{X}$ be a reordering of X, such that the creasing along X changes it sign and the right/left position of the faces is swapped. Suppose that f contains the face left of $\bar{X}$ and right of X, and that g contains the other face incident to X. Then we have by Lemma 2

$$\omega_{\bar{X}} = c(f, g, \bar{X}) = -c(f, g, X) = c(g, f, X) = \omega_X.$$

$\square$

We remark that the stress that we define is an equilibrium stress in the classical sense. In fact, in three dimensions it corresponds to the stress that is given by the Maxwell–Cremona correspondence. In particular, the formulation based on (1) can be found in a similar form in Hopcroft and Kahn [HK92, Equation (11)].

2.5 Convexity

The difficult part for the height assignment is to choose the heights such that the resulting lifting gives a *convex* realization of the simplicial complex. To guarantee the convexity of the final embedding we use the stresses induced by the lifting as a certificate. By knowing the signs of all stresses we can determine if the selected z-coordinates produce a convex embedding with help of the following lemma.

Lemma 4 *Let P be the lifting of a flat embedding, such that all vertices have been assigned with a nonnegative z-coordinate. If*

- *for all f_X incident to f_B we have $\omega_X < 0$ and*
- *for all other f_X we have $\omega_X > 0$,*

then P is a convex polytope.

Proof. As a preliminary step we show that P is *locally convex*. By this we mean that for every facet $f_T \neq f_B$ of P spanned by T, all adjacent facets lie "below" the hyperplane spanned by T. Let f_S be a face that is adjacent to f_T along the ridge f_X. The sequences S and T are ordered such that they coincide on the first $d-2$ elements, that is X. We assume for now that f_T is right of f_X, which implies $[\![T]\!] < 0$.

Case 1 ($S \neq B$): We denote by $\mathbf{r}$ the vertex in S that is not in T projected into the $z = 0$ hyperplane. By definition f_S is left of f_X and hence $[\![S]\!] > 0$. Since $\omega_X > 0$, we know that $c(f_S, f_T, X) > 0$. This leads to

$$\omega_X > 0 \iff c(f_S, f_T, X) > 0 \iff \frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]} > 0 \iff z_T(\mathbf{r}) > z_S(\mathbf{r}).$$

Hence the interior of f_S lies below $h(T)$.

Case 2 ($S = B$): The only difference here is that by assumption $\omega_X < 0$ and that $[\![S]\!]$ and $[\![T]\!]$ have the same sign, i.e., $[\![S]\!] < 0$. Therefore,

$$\omega_X < 0 \iff c(f_S, f_T, X) < 0 \iff \frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]} < 0 \iff z_T(\mathbf{r}) > z_S(\mathbf{r}).$$

Again, the interior of f_S lies below $h(T)$.

For both cases we have assumed that f_T is right of f_X. If f_T would be left of f_X instead, then the sign of $[\![S]\!]$ would change and furthermore, $\omega_X = c(f_T, f_S, X) = -c(f_S, f_T, X)$ by Lemma 2. Both effects cancel, i.e., for case 1 we have

$$\omega_X > 0 \iff c(f_S, f_T, X) < 0 \iff \frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]} < 0 \iff z_T(\mathbf{r}) > z_S(\mathbf{r}),$$

and the same "cancellation" happens in case 2.

We have shown that P is locally convex and will extend this to global convexity. This follows in our setting by Theorem 2.3.20 from the book of De Loera, Rambau and Santos [DLRS10]. For completeness we also give the detailed argument in this place. In particular, we show that for every facet $f_T \neq f_B$ all other vertices of P lie below $h(T)$. We first generalize the observations that proved the local convexity. In the above estimations we can pick as $\mathbf{r}$ any vertex that lies in the halfspace of $\mathbf{R}^{k-1}$ that is bounded by the supporting hyperplane of $\pi(f_X)$ and that does not contain $\pi(T)$. It still holds that $z_T(\mathbf{r}) > z_S(\mathbf{r})$.

We are now ready to prove that P is a convex polytope. Let $\mathbf{p}$ be a vertex of P not in T with z-coordinate z_p and let $\pi(\mathbf{p}) = \mathbf{r}$. Consider a segment ℓ in $\mathbf{R}^{k-1}$ that connects $\mathbf{r}$ with a point on $\pi(T)$. We can assume that ℓ avoids the vertices of $\pi(P)$ in its interior, otherwise we perturb ℓ slightly. Let the supported hyperplanes of the facets visited when traversing ℓ be (in order) $h(T) = h(1), h(2), \dots h(k)$. Due to the local convexity we have for $\mathbf{r} = \pi(\mathbf{p})$ that $z_{h(k)}(\mathbf{r}) < z_{h(k-1)}(\mathbf{r})$. Furthermore, due to the generalized observation in the previous paragraph we have for all $j \leq k$ that $z_{h(j)}(\mathbf{r}) < z_{h(j-1)}(\mathbf{r})$. This yields

$$z_T(\mathbf{r}) = z_{h(1)}(\mathbf{r}) > z_{h(2)}(\mathbf{r}) > \cdots > z_{h(k)}(\mathbf{r}) = z_p,$$

which proves the assertion for $f_T \neq f_B$. The base face f_B lies in the $z = 0$ hyperplane, which is clearly a bounding hyperplane since all vertices have positive z-coordinates. Thus all supporting hyperplanes of the facets of P are bounding hyperplanes and therefore P is convex. $\square$

2.6 Keeping track of the stresses while stacking

We discuss next how a stacking operation changes the associated stresses in a lifted barycentric subdivision. Let Δ_D be a $(d-1)$-simplex of a simplicial complex embedded in $\mathbf{R}^d$. We stack the new vertex $\mathbf{p}$ on top of D. To determine the stacking geometrically, we describe the location of $\mathbf{r} := \pi(\mathbf{p})$ inside $\pi(\Delta_D)$ with barycentric coordinates, that is we specify the absolute areas of the $(d-1)$-simplices containing $\mathbf{r}$ in the projection. Additionally, we describe how far the z-coordinate z_p of $\mathbf{p}$ lies above $h(D)$. This will be denoted by $\zeta := z_p - z_D(\mathbf{r})$. We refer to ζ as the *vertical shift*. For convenience we assume that the stacking process increases the z-coordinate of the stacked point, since this will always be the case in the following. However, the following observations can be easily generalized for stacking operations that decrease the z-coordinates of the stacked point.

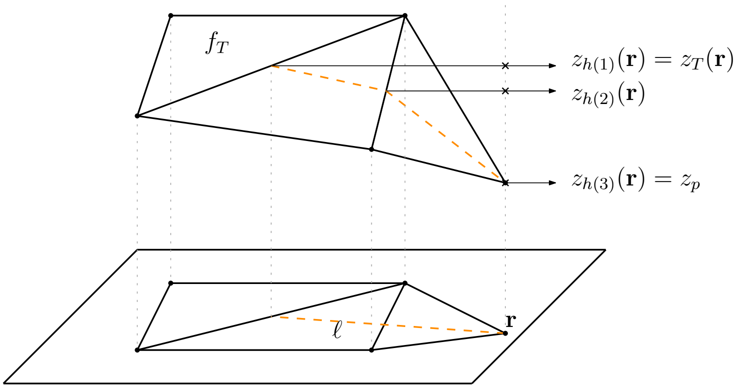

Figure 3: Proof of Lemma 4: A cascading sequence of inequalities (coming from the local convexity of P) shows that every vertex of P lies below the supporting hyperplane of a facet $f_T \neq f_B$.

The addition of $\mathbf{p}$ has two effects. First of all, by stacking $\mathbf{p}$ we create new ridges, furthermore, the existing ridges at the boundary f_D are now incident to new facets and hence the corresponding stress is altered. We refer to the newly introduced ridges as the *interior ridges*, and to the ridges on the boundary of f_D as the *exterior ridges* of the stacking operation.

We discuss first how the stresses on the exterior ridge are altered.

Lemma 5 *Assume that we have stacked a new vertex* $\mathbf{p}$ *on a convex polytope realized in* $\mathbf{R}^d$ *with stress* ω *as described in Lemma 4. Let* X *be the point set of an exterior ridge* f_X *(as defined above). After the stacking the new stress* $\hat{\omega}$ *equals*

$$\hat{\omega}_X := \omega_X - \frac{\zeta}{|[\![S]\!]|},$$

where $S = X \circ \mathbf{p}$ *and* ζ *is the vertical shift of the stacking.*

Proof. Let T be the ordered point set, such that f_T and f_S are adjacent after stacking $\mathbf{p}$, and f_X is the intersection of f_S and f_T. We denote as h the hyperplane that contains the face f_D on which $\mathbf{p}$ was stacked onto; see Fig. 4.
Assume for now that f_S is left of f_X, which implies that $[\![S]\!] > 0$. This gives for $\mathbf{r} = \pi(\mathbf{p})$

$$\hat{\omega}_X = c(f_S, f_T, X) = \frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]} = \frac{z_T(\mathbf{r}) - (z_h(\mathbf{r}) + \zeta)}{[\![S]\!]} = \omega_X - \frac{\zeta}{[\![S]\!]} = \omega_X - \frac{\zeta}{|[\![S]\!]|}.$$

If f_S is however right of f_X we obtain

$$\hat{\omega}_X = -c(f_S, f_T, X) = -\frac{z_T(\mathbf{r}) - z_S(\mathbf{r})}{[\![S]\!]} = -\frac{z_T(\mathbf{r}) - (z_h(\mathbf{r}) + \zeta)}{[\![S]\!]} = \omega_X + \frac{\zeta}{[\![S]\!]} = \omega_X - \frac{\zeta}{|[\![S]\!]|},$$

since $[\![S]\!] < 0$. $\square$

Lemma 6 *Assume we have stacked a vertex* $\mathbf{p}$ *on top of some face* f_D*. Let* X *be the point set of an interior ridge* f_X *as defined above. Furthermore, let* $S := X \circ \mathbf{s}_d$ *and* $T := X \circ \mathbf{t}_d$ *be the sequence of vertices of the faces separated by* X*. After the stacking the stress on* X *equals the positive number*

$$\omega_X := \zeta \left| \frac{[\![D]\!]}{[\![S]\!][\![T]\!]} \right|.$$

Proof. Fig. 5 depicts the situation described in the lemma. The facet f_S might be either left or right of f_X. Due to Equation (3) we have

$$\omega_X = \pm c(f_S, f_T, X) = \pm \frac{[T \circ \mathbf{s}_d]}{[\![T]\!][\![S]\!]}.$$

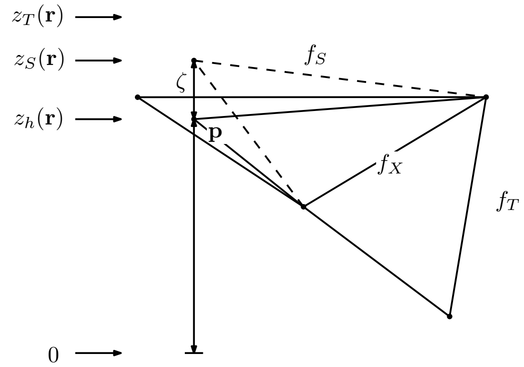

Figure 4: The situation (for $d = 3$) as discussed in Lemma 5: the effect of stacking the vertex $\mathbf{p}$ for an exterior ridge.

Assume we have doubled f_D and stack with an ε-small ζ at the "top side" of the induced complex. This will surely generate a convex d-simplex, and by the computations in the proof of Lemma 4 the sign of the induced ω_X is positive. Clearly, the sign of ω_X does not change if we increase the vertical shift.

We know that $|[T \circ \mathbf{s}_d]|$ is $d!$ times the volume of the corresponding simplex. This volume on the other hand can be expressed as $\zeta \cdot |vol(\pi(\Delta_D))|/d$, and $|[\![D]\!]|$ equals $(d-1)!$ times the volume of $\pi(\Delta_D)$. Hence we have $|[T \circ \mathbf{s}_d]| = \zeta \cdot |[\![D]\!]|$ and the statement of the lemma follows. $\square$

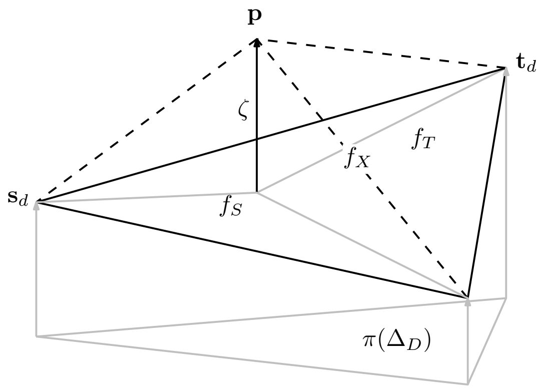

Figure 5: The situation (for $d = 3$) for an interior ridge as discussed in Lemma 6.

2.7 The tree-representation of a stacked polytope

The combinatorial structure of a polytope is typically encoded by its face lattice. If the polytope is a stacked polytope, we can also describe its structure by "recording" in which way we have carried out the stacking operations. The most natural way to keep track of the stacking operations is an ordered rooted tree. Let P be a stacked polytope, then the *tree-representation* of P is a d-ary tree $\mathcal{T}(P)$, which is defined as follows: The leaves of the trees are in one-to-one correspondence to the facets of P, with the exception of the face f_B, which is not present in the tree. Interior nodes correspond to d-cliques in the graph of P that used to be facets at some point during the stacking process.[3] The root represents the initial copy of f_B in the beginning. The children of a node v represent the faces that were introduced by stacking a vertex onto the face associated with v.

[3] When considering an intermediate configuration in the stacking process, we refer to the d-cliques of P that are faces in the intermediate polytope also as *facets*, although they are not necessarily facets of P.

To reproduce the combinatorial structure from $\mathcal{T}(P)$ we fix an ordering for the edges emanating from an interior node v, such that it is possible to map the children of v to the faces generated by stacking onto v's face in a unique way. By this the mapping between the faces of P and the nodes of $\mathcal{T}(P)$ can be reproduced. If P has n vertices, then $\mathcal{T}(P)$ has $(n-d)(d-1)+1$ leaves and $n-d$ interior nodes.

We will use the tree-representation of P to specify the geometry of the flat embedding of P. This can be achieved by assigning for every leaf v of $\mathcal{T}(P)$ a rational number (weight), which corresponds to the volume of v's face in the flat embedding of P. More precisely, the weight specifies the volume of v's face times $(d-1)!$. To emphasize this relationship, we call the weights *face-weights*. We extend the face-weight assignments for the interior nodes by summing for every node the face-weights of its children recursively up. After we have determined the location of f_B, such that its volume is $(d-1)!$ times the face-weight of the root of $\mathcal{T}(P)$, the coordinates of the whole flat embedding are determined. In particular, for every stacking operation, the (normalized) face-weights specify the location of the stacked vertex since they denote its barycentric coordinates. Thus, by traversing $\mathcal{T}(P)$ we can produce the flat embedding incrementally in a unique way.

3 The embedding algorithm

We assume that the combinatorial structure of the stacked polytope is given in form of its tree-representation $\mathcal{T}(P)$. The embedding algorithm works in three steps. First we generate the face-weights for $\mathcal{T}(P)$ and fix the coordinates for the face f_B. This will give us a flat embedding of the polytope. In the next step we "lift" the polytope, by defining for every vertex v_i the *vertical shift* ζ_i. Assume that we have stacked v_i onto the face f. By construction we pick always positive vertical shifts. By carefully choosing the right vertical shifts we obtain an embedding of P as a *convex* polytope, however this embedding does not fit on a polynomial grid. In the final step we round the points to appropriate grid points, while maintaining convexity.

3.1 Balancing face-weights

We apply a technique from data structure analysis called the *heavy path decomposition* (see Tarjan [Tar83]). Roughly speaking, it decomposes a tree into paths, such that the induced hierarchical structure of the decomposition is balanced. We continue with a brief review of the heavy path decomposition. Let u be a non-leaf of a rooted tree T with root r ($r = u$ is possible). We denote by T_u the subtree of T rooted at u. Let v be the child of u such that T_v has the largest number of nodes (compared to the subtrees of the other children of u), breaking ties arbitrarily. We call the edge (u, v) a *heavy edge*, and the edges to the other children of u *light edges*. The heavy edges induce a decomposition of T into paths, called *heavy paths*, and light edges; see Fig. 6. The node on a heavy path that is closest to the root is called its *top node*. We call a heavy path with its incident light edges (ignoring the possible edge from its top node to its parent) a *heavy caterpillar*. The heavy path decomposition decomposes the edges of T into heavy caterpillars. We say that two heavy caterpillars are adjacent if their graphs would be adjacent subgraphs in T. This adjacency relation induces a hierarchy, which we represent as a rooted tree $\mathcal{H}(T)$. The nodes in $\mathcal{H}(T)$ are the heavy caterpillars, and its edges represent the adjacency relation between caterpillars. The root of $\mathcal{H}(T)$ is the heavy caterpillar that includes the root of T. When (u, v) is a light edge and u is the parent of v then the size of T_v is at most half as big as the size of T_u. Hence, every root-leaf path in T can visit at most $\log t$ light edges for t being the size of T. Fig. 6 shows an example of a tree-representation and its associated hierarchy as a tree.

The assignment of the face-weights is guided by the heavy path decomposition of $\mathcal{T}(P)$. We call a node v *balanced* if (1) the face-weights of v's light-edge children are all the same and (2) the face-weight of the heavy edge child is not smaller than the face-weight of every light edge child. If every node is balanced we call the face-weight assignment balanced. An example of a balanced tree-representation is depicted in Fig. 8.

Lemma 7 *We can find a balanced set of integer face-weights for $\mathcal{T}(P)$ such that the face-weight associated with the root is at most $dn^{\log 2d}$ and no face-weight is less than 1. Here, n denotes the number of vertices in the corresponding graph of the stacked polytope.*

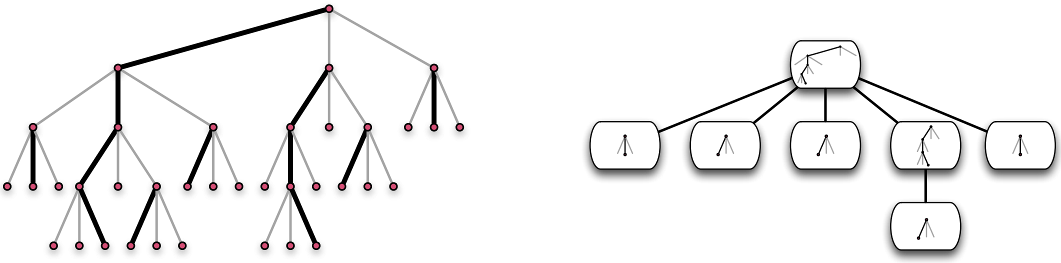

(a) (b)

Figure 6: A tree-representation $\mathcal{T}$ of a stacked 3-polytope (a) and the corresponding hierarchy based on the caterpillars as tree $\mathcal{H}(\mathcal{T})$ (b).

Proof. Let $H = \mathcal{H}(T)$ be the tree representing the hierarchy of the heavy caterpillars for some d-ary tree T. The height of H is the length (number of edges) of its longest root-leaf path. For T we call the height of $\mathcal{H}(T)$ its *hpd-height*.

To prove the lemma we use the following claim. Since the proof of the claim will be constructive it will also induce a strategy of how to set the face-weights.

Claim: For all $k \geq 0$ the following holds: If T is a d-ary tree with m leaves and hpd-height k then we can find a set of balanced face-weights, such that the face-weight of the root is at most md^k.

We prove the claim by induction on k. A tree with hpd-height 0 is a heavy caterpillar, whose subtrees on the light edges are leaves. In this case it suffices to set the face-weights of all leaves to 1 and then propagate the face-weights to the interior nodes. Clearly, this face-weight assignment is balanced and the face-weight of the root equals the number of leaves. Thus, the claim holds for this (base) case.

We continue with the induction step. Assume that the claim holds for all trees with hpd-height less then k. We can now balance the face-weights of any tree of hpd-depth k as follows (see also Fig. 7 for an illustration): Let the heavy path incident to the root be h. The subtrees connected to h via a light edge have all depth less than k. For these subtrees we use our induction hypothesis and (recursively) determine the corresponding face weights and propagate these weights to h. If a subtree is a single leaf it will be assigned with a face-weight of 1. We are left with balancing the nodes of h. So let v be a node on h and let u_v^+ be one of its light edge children with the largest face-weight. We will now make each light child $u \neq u_v^+$ as "heavy" as u_v^+. To do so, we determine the difference δ between the face-weight of u_v^+ and the face-weight of u. Next, we increase the face-weight of every node on the heavy path with top node u by δ. Notice that this keeps the face-weights in the subtree rooted at u balanced. To make the updates consistent we also increase the face-weight on h from v to its root by δ. We continue with the other children of u in the same fashion and then repeat this for every node on h. We also have to guarantee that the face-weight of every light edge child is not larger than the face-weight of its heavy edge child. To ensure this, we increase the face-weight of the leaf of h by δ_r and propagate this weight along h, where δ_r is the largest face-weight of one of the light edge subtrees hanging off of h. Note that this is done only once for the heavy path h. By this we have obtained a balanced set of face-weights for T.

We are left with bounding the new face-weight of the root. Let S be the sum of the face-weights of all the light edge subtrees of h before balancing the nodes of h, which is also the face-weight of the root minus 1 at this time. By the induction hypothesis we have $S \leq md^{k-1}$. When balancing the light subtrees on a vertex v we increased the face-weight of the root by some δ at most $d-2$ times. The "charge" δ was not larger than the face-weight of the corresponding u_v^+. The total increment at this stage is therefore less than $(d-2)S$. The final increment by δ_r is at most $S-1$. Therefore, we have that the face-weight of the root is at most

$$(S+1) + (d-2)S + (S-1) = dS \leq dm^k$$

and the proof of the claim follows.

The statement of the lemma follows now from the claim. According to Subsec. 2.7 the tree $\mathcal{T}(P)$ has less than $(d-1)n$ leaves. When traversing a light edge (u, v) the number of interior nodes is at least halved. The tree $\mathcal{T}(P)$ has less than n interior nodes. Thus, every path in $\mathcal{T}(P)$ from the top node of a heavy path to the root of $\mathcal{T}(P)$ visits at most $\log n$ light edges. As a consequence, the maximal hpd-height is bounded by $\log n$. The total face-weight of the root is therefore at most $(d-1)n \cdot d^{\log n} < dn \cdot n^{\log d} = dn^{\log 2d}$. $\square$

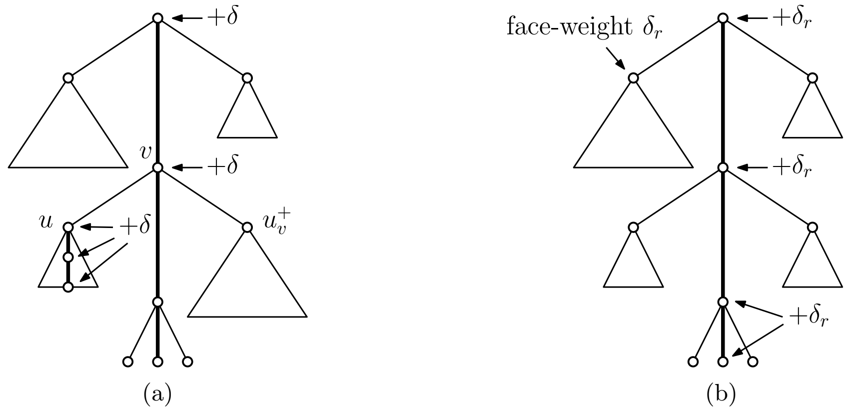

Figure 7: Modifications of the face-weight as done in the proof of Lemma 7. (a) Sketch how to make all light children have the same face-weight. (b) Sketch how to keep the heavy edge child the child with largest face-weight.

In the following we denote the face-weight of the root as R. To finish the definition of the flat embedding we fix the shape of f_B as follows: One of the vertices of B lies at the origin. We set $L = \sqrt[d-1]{R}$ and place the other vertices of B at $L \cdot \mathbf{e}_i$, for $\mathbf{e}_i$ being a vector of the standard basis of $\mathbf{R}^{d-1}$, such that B spans the $(d-1)$-simplex Δ_B. By construction the volume of the simplex equals $1/(d-1)!$ times the face-weight of the root of $\mathcal{T}(P)$. Without loss of generality we also assume for the remainder that $[\![B]\!] > 0$. It follows that

$$[\![B]\!] = R \leq dn^{\log(2d)}.$$

By our choice of f_B the volume of every face in the flat embedding equals $(d-1)!$ times the volume of its face-weight.

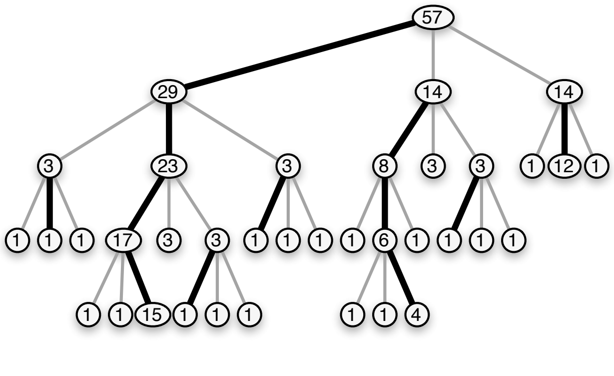

Figure 8: The tree-representation of Fig. 6 with balanced face-weights, denoted in the nodes.

3.2 Assigning heights

After we have assigned the face-weights we specify the heights of the lifting by determining the vertical shift ζ_i for every vertex v_i based on the face-weights. Let us assume vertex v_i was stacked as $\mathbf{p}$ onto some face f_D. Let the boundary of f_D be formed by the ridges $f_{X_1}, \ldots, f_{X_d}$. The stacking introduces the new facets $f_{Y_1}, \ldots, f_{Y_d}$, where for $1 \le j \le d$ we have $Y_j := X_j \circ \mathbf{p}$. The nodes of the faces f_{Y_j} in $\mathcal{T}(P)$ that are connected to the node of f_D via a light edge have all the same face-weight. Let this weight be B_i. The only (possible) other face-weight for a node associated with one of the faces f_{Y_j} is denoted by A_i. We set as the vertical shift for the vertex v_i

$$\zeta_i := A_i \cdot B_i. \tag{5}$$

Note that $A_i \ge B_i$, since the face-weights were balanced.

Lemma 8 *The embedding induced by the face-weights and the vertical shifts ζ_i guarantees:*

1.) For every ridge f_X not on f_B we have $\omega_X \ge 1$.

2.) For every ridge f_X on f_B we have $0 > \omega_X \ge -R \ge -d \cdot n^{\log(2d)}$.

Proof. We first bound the stresses on faces not on f_B. Recall that by construction the face-weight of a facet f_D coincides with $|[\![D]\!]|$. We study how the stress on ω_X evolves during the stacking process. The initial value of ω_X is assigned by some stacking operation that introduced f_X. We assume that this stacking operation stacked a vertex on the face f_D. The face-weights of the new facets introduced by the stacking are all the same, namely B_i, except for one possible larger face-weight, namely A_i. Let f_S and f_T be the two facets incident to f_X such that $|[\![S]\!]| \ge |[\![T]\!]|$. By Lemma 6, we have that at this moment

$$\omega_X = \zeta_i \left| \frac{[\![D]\!]}{[\![S]\!][\![T]\!]} \right| = A_i B_i \left| \frac{[\![D]\!]}{[\![S]\!][\![T]\!]} \right| \ge |[\![D]\!]| \,,$$

since $|[\![S]\!]| \le A_i$ and $|[\![T]\!]| = B_i$.

This positive initial stress decreases when stacking on a facet that has f_X on the boundary. So assume we stacked $\mathbf{p}_k$ on such a facet. Let $C_X(k)$ denote the amount of the decrement due to this stacking. By Lemma 5 we have $C_X(k) = \zeta_k / |[\![S_k]\!]|$, where $S_k = X \circ \mathbf{p}_k$. Recall that $\zeta_k = A_k B_k$, where A_k and B_k are the two different values of face-weights of faces introduced by stacking $\mathbf{p}_k$. It follows that $|[\![S_k]\!]| \in \{A_k, B_k\}$, and therefore $C_X(k) = \{A_k, B_k\} \setminus \{|[\![S]\!]|\}$ equals either A_k or B_k, and hence, is an integer. We charge the value of $C_X(k)$ to a face-weight of a face $f_{Y_k} \ne f_{S_k}$ that was introduced when stacking $\mathbf{p}_k$. Stacking operations that decrease ω_X further, stack either onto f_{S_k}, or onto the opposite facet incident to f_X (see Fig. 9). As a consequence, the projected facets $\{\pi(f_{Y_j}) \mid \text{stacking of } \mathbf{p}_j \text{ decrements } \omega_X\}$ have disjoint interiors and are properly contained inside $\pi(\Delta_D)$. Since it is not possible to cover $\pi(\Delta_D)$ with these faces completely, the difference $|[\![D]\!]| - \sum_j C_X(j)$ is at least 1 and therefore $\omega_X \ge |[\![D]\!]| - \sum_j C_X(j) \ge 1$.

For the stresses on the boundary we can argue as follows. The stress on X is a combination of several negative charges, but without having an initial positive charge. By the above argument, the negative charges are attributed to faces with disjoint interiors. All these faces are contained inside $\pi(\Delta_B)$ and hence $|\omega_X|$ is at most $[\![B]\!]$. By Lemma 7 we have that $[\![B]\!] = R \le d \cdot n^{\log(2d)}$, and the lemma follows. □

3.3 Rounding to grid points

Let us wrap up what we have constructed so far. Based on the tree-representation we have defined weights for each face. This gave rise to a realization of a projection of the stacked polytope. As a next step we have computed a lifting based on the face-weights. This construction produces a convex realization of the desired stacked polytope, we even know that the stresses are polynomially related. However, the realization does not lie on a polynomial grid yet. To obtain an integer realization we round the coordinates of the points down. The rounding will be carried out in two steps. First we perturb all coordinates such that they are multiples of some parameter α,

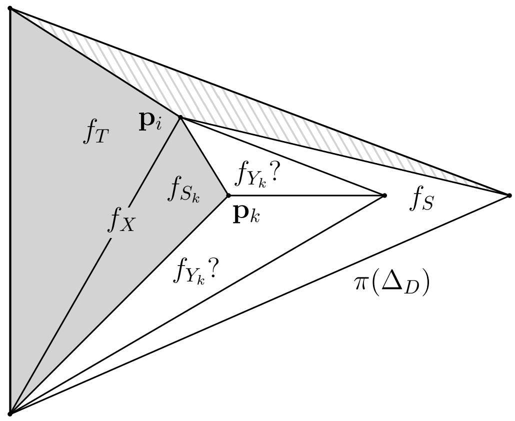

Figure 9: When stacking $\mathbf{p}_k$ the stress on X is decreased by at most the face-weight charged to f_{Y_k}. There are two candidates for f_{Y_k}. All further negative charges will be attributed to "volumes" contained inside S_k and T. The face with the crossed-line pattern will never be charged to ω_X.

resp. α_z for the z-coordinates. In a second step we scale the perturbed embedding by multiplying all coordinates with $1/\alpha$, resp. with $1/\alpha_z$, for the z-coordinates. The details for the rounding procedure are little bit more subtle: We start with rounding the coordinates in the flat embedding, then we update the vertical shifts ζ_i slightly. The z-coordinates are rounded with respect to the lifting defined by the modified vertical shifts.

Since projected volumes play an important role in our construction we discuss as a first step how the rounding will effect these volumes.

Lemma 9 *If we round the coordinates of the flat embedding down, such that every coordinate is a multiple of α we have that for every facet f_X*

$$|[\![X]\!] - [\![X']\!]| \le \alpha d^2 L^{d-2},$$

where X' denotes the points X after the rounding.

Proof. Remember that the points X are contained inside the simplex Δ_B, which is spanned by the standard basis vectors of $\mathbf{R}^{d-1}$ scaled by $L = \sqrt[d-1]{R}$ and the origin. We first show how the rounding of the first coordinate for every $x \in X$ effects $[\![X]\!]$. The point set X after rounding the first coordinate is named X^1. Let $E(1)$ denote the change of $[\![X]\!]$ in the worst case, that is, $E(1) = \max_X \left|[\![X]\!] - [\![X^1]\!]\right|$. We denote by $x_i \in \mathbf{R}$ the first coordinate (the x-coordinates) of the ith point in X, and for a point sequence X we denote by X_{-i} the set with removed x-coordinates and without the ith point of X. Note that $X_{-i} = X^1_{-i}$. Let the change of the coordinates x_i be $\varepsilon_i \le \alpha$. We estimate $E(1)$ using the Laplace expansion for the determinants (along the row for the x-coordinates) by

$$\begin{aligned}
E(1) = \left|[\![X]\!] - [\![X^1]\!]\right| &\le \left|\sum_{i=1}^{d}(-1)^i x_i [\![X_{-i}]\!] - \sum_{i=1}^{d}(-1)^i (x_i+\varepsilon_i)[\![X^1_{-i}]\!]\right| \\
&\le \left|\sum_{i=1}^{d}(-1)^i x_i [\![X_{-i}]\!] - \sum_{i=1}^{d}(-1)^i (x_i+\varepsilon_i)[\![X_{-i}]\!]\right| \\
&\le \sum_{i=1}^{d} |\alpha [\![X_{-i}]\!]| \\
&\le d\alpha L^{d-2}.
\end{aligned}$$

For the last transition we upper bounded the determinants $[\![X_{-i}]\!]$ by L^{d-2}. This bound follows, since the maximal volume of a $(d-2)$-simplex in Δ_B is spanned by the $d-1$ points of $B \setminus \mathbf{0}$. This point set would also maximize $[\![X_{-i}]\!]$ and hence this determinant is at most L^{d-2}.

Note that there was nothing special in rounding the *first* coordinate (while keeping the others fixed). If we would round only the kth coordinate, we obtain the same estimation. Thus, we can apply the rounding of a single coordinate one by one for each of the $(d-1)$-coordinates of $\mathbf{R}^{d-1}$. Every time we round in one coordinate we introduce an additive "error" of $d\alpha L^{d-2}$ to $|[\![X]\!] - [\![X']\!]|$. Thus in total we get the asserted bound of $|[\![X]\!] - [\![X']\!]| \leq \alpha d^2 L^{d-2}$. □

In the following we use $[\![X']\!]$ and similar expressions to denote the corresponding determinants after rounding.

Corollary 1 *If we round the coordinates of the flat embedding down, such that every coordinate is a multiple of α, we have for every facet f_X in the flat embedding*

$$1 - \alpha d^2 L^{d-2} \leq \frac{[\![X']\!]}{[\![X]\!]} \leq 1 + \alpha d^2 L^{d-2}.$$

Proof. The statement follows from Lemma 9 and the observation that for any facet f_X in the flat embedding we have $1 \leq |[\![X]\!]|$. □

As a next step we discuss, how to set the parameter α, such that the perturbed flat embedding will still have a positive stress. To describe the lifting (stress) we defined for every vertex in the flat embedding a vertical shift ζ_i, as given in (5). The definition of ζ_i was based on the face-weights obtained from the balanced tree-representation. We adjust the vertical shifts after the perturbation slightly. When stacking $\mathbf{p}_i$ we introduced d new faces. Let the faces f_{A_i} and f_{B_i} be the two faces out of the d newly introduced faces with the largest volume. We define

$$\zeta_i' := |[\![A_i']\!][\![B_i']\!]|.$$

Lemma 10 *When we pick as the perturbation parameter $\alpha = 1/(10d^2L^{d-2}R)$ then the perturbed flat embedding with the vertical shifts ζ_i' induces an embedding, whose interior stresses are at least $4/5$, and whose boundary stresses are negative and larger than $-2R$.*

Proof. We mimic the strategy of the proof of Lemma 8. Again, all faces in this proof are considered as projected into the $z = 0$ hyperplane. For the proof of the lemma the sign of the determinants $[\![\cdot]\!]$ is not important. For the sake of a simple presentation we misuse notation and simply write $[\![\cdot]\!]$ instead of $|[\![\cdot]\!]|$ in this proof.

Every stress ω_X is a combination of a positive stress and several negative stresses attributed to different stacking operations. Let ω_X^+ be the positive stress, that is the initial nonzero stress with respect to the stacking sequence, and let ω_X^- the absolute value of the sum of all negative stresses, such that $\omega_X = \omega_X^+ - \omega_X^-$. To bound ω_X we derive bounds for ω_X^+ and ω_X^-. We start with the bound on the positive stress. Assume ω_X^+ was introduced by stacking $\mathbf{p}_i$ at some face f_D, such that due to Lemma 6 we have

$$\omega_X^+ = \zeta_i' \left| \frac{[\![D']\!]}{[\![S']\!][\![T']\!]} \right|.$$

Here, f_S and f_T are the two faces introduced by stacking $\mathbf{p}_i$ that contain f_X. The height ζ_i' is defined as the product of two face-weights corresponding to f_{A_i} and f_{B_i}. Assume that $[\![A_i']\!] \geq [\![B_i']\!]$ and that $[\![S']\!] \geq [\![T']\!]$. By the definition of ζ_i' we have that $[\![A_i']\!] \geq [\![S']\!]$ and $[\![B_i']\!] \geq [\![T']\!]$. Therefore, $\omega_X \geq [\![D']\!]$.

The value of ω_X^- is composed of several "charges". Whenever we stack inside a face that contains f_X we increase ω_X^-. Let us study now one of these situations. Assume we stack $\mathbf{p}_k$ inside a face that contains f_X. Let f_{S_k} be the new face that contains f_X. By Lemma 5 we increase ω_X^- by $|\zeta_j'/[\![S_k']\!]| := inc_{X,k}$. Due to the balanced face-weights we had in the unperturbed setting only two different new "face volume values" when stacking a vertex. Hence for $\zeta_i' := |[\![A_k']\!][\![B_k']\!]|$, we had either $[\![S_k]\!] = [\![A_k]\!]$, or $[\![S_k]\!] = [\![B_k]\!]$. We define

$$C_k := \begin{cases} B_k & \text{if } [\![S_k]\!] = [\![A_k]\!] \\ A_k & \text{if } [\![S_k]\!] = [\![B_k]\!] \end{cases}$$

and set $\delta_+ := 1+\alpha d^2 L^{d-2}$ and $\delta_- = 1-\alpha d^2 L^{d-2}$. If $[\![S_k]\!] = [\![A_k]\!]$ we have according to Corollary 1

$$inc_{X,k} = \left| \frac{[\![A'_k]\!]}{[\![S'_k]\!]} [\![B'_k]\!] \right| \leq \left| \frac{\delta_+ [\![A_k]\!]}{\delta_- [\![S_k]\!]} [\![B'_k]\!] \right| = \frac{\delta_+}{\delta_-} [\![B'_k]\!] = \frac{\delta_+}{\delta_-} [\![C'_k]\!].$$

The remaining case $[\![S_k]\!] = [\![B_k]\!]$ is completely symmetric and does also give $inc_{X,k} \leq \frac{\delta_+}{\delta_-} [\![C'_k]\!]$. Let K be the set of vertex indices, whose stacking contributed to ω_X^-. As noticed in Lemma 8 (see also Fig. 9), for any two distinct $s, t \in K$ we have that f_{C_s} and f_{C_t} have disjoint interiors, and furthermore the set $\bigcup_{k \in K} f_{C_k}$ is contained in the perturbed simplex $\pi(\Delta_D)$, but "misses" at least one face. By Corollary 1 the face-weight of the missing face is at least δ_-. Therefore,

$$\omega_X^- = \sum_{k \in K} inc_{X,k} \leq \sum_{k \in K} \frac{\delta_+}{\delta_-} [\![C'_k]\!] = \frac{\delta_+}{\delta_-} \sum_{k \in K} [\![C'_k]\!] \leq \frac{\delta_+}{\delta_-} \left([\![D']\!] - \delta_- \right) = \frac{\delta_+}{\delta_-} [\![D']\!] - \delta_+.$$

We finish the proof by combining the bounds for ω_X^+ and ω_X^-. When we pick $\alpha = 1/(10d^2 L^{d-1} R)$ as specified in the lemma then we get $\delta_+ = 1 + \frac{1}{10R}$ and $\delta_- = 1 - \frac{1}{10R}$. We can now obtain the following bound for ω_X when X is an interior ridge (note that $\pi(\Delta_{D'}) \subseteq \pi(\Delta_B)$) :

$$\begin{aligned}
\omega_X &\geq \omega_X^+ - \omega_X^- \\
&\geq [\![D']\!] - \left(\frac{\delta_+}{\delta_-} [\![D']\!] - \delta_+ \right) \\
&= \delta_+ - [\![D']\!] \left(\frac{\delta_+}{\delta_-} - 1 \right) \\
&\geq \delta_+ - [\![B]\!] \left(\frac{\delta_+}{\delta_-} - 1 \right) = \delta_+ - \delta_+ R \left(\frac{\delta_+}{\delta_-} - 1 \right) \\
&= \frac{80R^2 - 2R - 1}{100R^2 - 10R} \\
&\geq \frac{4}{5}
\end{aligned}$$

If X is a boundary face we have $\omega_X^+ = 0$. Notice that for our choice of α we have $\delta_+/\delta_- \leq 2$, for $R \geq 3$. We conclude that for a boundary face we have

$$\omega_X \geq -\left(\frac{\delta_+}{\delta_-} [\![B']\!] + \delta_+ \right) \geq -\frac{\delta_+}{\delta_-} [\![B']\!] \geq -2[\![B']\!] \geq -2R.$$

□

The choice of the parameter $\alpha = 1/(10d^2 L^{d-2} R)$ ensures, that no volume of a face will flip its sign. In particular, by Lemma 9, the change of volume is less than $1/(10[\![B]\!])$.

To construct an integer realization we need to round the z-coordinates as well. We round the z-coordinates such that every z-coordinate is a multiple of α_z, where α_z is some value to be determined later. The final analysis requires an upper bound for the maximal z-coordinate before rounding, which we give in the following lemma.

Lemma 11 *For z_{max} being the maximal z-coordinate in the lifting of the perturbed flat embedding induced by the vertical shifts ζ'_i we have that*

$$0 < z_{max} < 2R^2.$$

Proof. The z-coordinates of all vertices not on f_B are positive since all vertical shifts are positive. We consider the perturbed flat embedding. Take any boundary ridge f_X of f_B. By Lemma 10 we have that $-\omega_X < 2R$. Let $\mathbf{p}_i$ be the vertex with the highest z-coordinate $z_{\max}$ and set $\mathbf{r} = \pi(\mathbf{p}_i)$. Now take the facet f_Y adjacent to f_B via f_X (we let Y and B coincide on the first $d-2$ vertices). Due to the convexity of the lifting f_Y supports a bounding hyperplane, and therefore $z_Y(\mathbf{r}) \geq z_{\max}$. By Equation (1) we have $\omega_X [\![Y]\!] = z_B(\mathbf{r}) - z_Y(\mathbf{r}) = -z_Y(\mathbf{r})$. Since f_Y is properly contained inside f_B it follows that

$$z_{\max} \leq z_Y(\mathbf{r}) = -\omega_X [\![Y]\!] < -\omega_X [\![B]\!] \leq 2R^2.$$

□

By rounding the z-coordinates we might violate the convexity of the lifting. The following lemma shows us how to carry out the rounding of the z-coordinates after we have rounded the coordinates of the flat embedding, such that the resulting embedding remains a convex polytope.

Lemma 12 *By setting $\alpha_z = 1/(3R)$ and rounding such that all z-coordinates are multiples of α_z the lifting defined by ζ_i' and the perturbed flat embedding will remain an embedding of a convex polytope.*

Proof. Let z_i' be the coordinate z_i after rounding and, similarly, let ω_X' be the stress after rounding the z-coordinates. We have that $0 \le z_i - z_i' \le \alpha_z$. By Lemma 10 the interior ridges in the perturbed flat embedding have a stress ω_X that is at least $4/5$. Let f_X be any interior ridge that is incident to the facets f_S and f_T. By Equation (3) we can express ω_X' in terms of $[\mathcal{T}']$, $[\![S']\!]$, and $[\![T']\!]$, where $\mathcal{T}'$ is formed by the point set $S' \cup T'$. The value of $[\mathcal{T}']$ can be expressed as $\sum_{i\in I} z_i'[\![A_i']\!]$, where I is the index set of the vertices of $\mathcal{T}'$ and A_i' is given by $\mathcal{T}' \setminus \{\mathbf{p}_i'\}$ in some appropriate order (Laplace expansion). We split the set I into $I_+ := \{i \in I \mid [\![A_i]\!] \ge 0\}$ and $I_- := I \setminus I_+$. As usual f_B denotes the boundary face. Note that the projections of the simplices spanned by the sets A_i' double-cover the projection of $\mathcal{T}'$ into the $z = 0$ hyperplane, and therefore $\sum_{i\in I} |[\![A_i']\!]| \le 2[\![B']\!] \le 2[\![B]\!] \le 2R$. Moreover, by Corollary 1 and our choice of α, we have $|[\![S']\!][\![T']\!]| \ge (1 - 1/(10R))^2$. The stress on an interior ridge f_X after rounding (ω_X') can be bounded as follows

$$\begin{aligned}
\omega_X' = \frac{[\mathcal{T}']}{[\![S']\!][\![T']\!]} = \frac{\sum_{i\in I} z_i'[\![A_i']\!]}{[\![S']\!][\![T']\!]} &= \frac{\sum_{i\in I_+} z_i'|[\![A_i']\!]|}{[\![S']\!][\![T']\!]} - \frac{\sum_{i\in I_-} z_i'|[\![A_i']\!]|}{[\![S']\!][\![T']\!]} \\
&\ge \frac{\sum_{i\in I_+} (z_i - \alpha_z)|[\![A_i']\!]|}{[\![S']\!][\![T']\!]} - \frac{\sum_{i\in I_-} z_i|[\![A_i']\!]|}{[\![S']\!][\![T']\!]} \\
&= \underbrace{\frac{\sum_{i\in I} z_i[\![A_i']\!]}{[\![S']\!][\![T']\!]}}_{t_1} - \underbrace{\frac{\sum_{i\in I_+} \alpha_z|[\![A_i']\!]|}{[\![S']\!][\![T']\!]}}_{t_2}
\end{aligned}$$

We observe that the $\sum_{i\in I} z_i[\![A_i']\!] = [\mathcal{T}_z']$, where $\mathcal{T}_z'$ denotes $\mathcal{T}$ after rounding in the flat embedding, but before rounding the z-coordinates. This gives us $t_1 \ge [\mathcal{T}_z']/([\![S']\!][\![T']\!]) = \omega_X \ge 4/5$. Moreover, as already noticed, $\sum_{i\in I_+} |[\![A_i']\!]| \le \sum_{i\in I} |[\![A_i']\!]| \le 2R$. After plugging in our choice for α_z and our bound for $[\![S']\!][\![T']\!]$ we obtain $t_2 \le 2R/(3R(1 - 1/(10R))^2)$. This yields

$$\omega_X' \ge t_1 - t_2 \ge \frac{4}{5} - \frac{2R}{3(1 - 1/(10R))^2 R} = \frac{4}{5} - \frac{200R^2}{3(10R-1)^2}.$$

Note that the last expression is a monotone increasing function which is positive for $R \ge 3$.

We are left with checking the sign for the stresses on the ridges that define the boundary of f_B. The stresses for these faces have to remain negative. Note that this is certainly the case, if all z-coordinates after the rounding remain positive. Before rounding the z-coordinates, every z-coordinate of a vertex not on f_B was at least as large as the smallest vertical shift ζ_i'. Since the vertical shifts are defined as the sum of two face-weights, we have that the nonzero z-coordinates are at least $(1 - 1/(10R))^2$. As observed earlier, rounding the z-coordinates decreases the z-coordinates by at most $1/(3R)$. Therefore we have for every $\mathbf{p}_i \notin B$

$$z_i' \ge \left(1 - \frac{1}{10R}\right)^2 - \frac{1}{3R} > 1 - \frac{1}{5R} - \frac{1}{3R} > 0,$$

for every $R \ge 3$. Hence, after rounding the z-coordinates the sign pattern of the stresses verifies the convexity of the perturbed realization. □

We now summarize our analysis and state the main theorem.

Theorem 1 *Every d-dimensional stacked polytope can be realized on an integer grid, such that all coordinates have size at most $10d^2R^2$, except for one axis, where the coordinates have size at most $6R^3$, for $R = dn^{\log(2d)}$.*

Proof. To get integer coordinates we multiply all coordinates after the rounding with $1/\alpha$, except the z-coordinates, which we multiply with $1/\alpha_z$. Since the maximal z-coordinate is by Lemma 11 at most $2R^2$, and we scale by $1/\alpha_z = 3R$ the bound for the z-coordinates in the theorem follows. All other coordinates are positive and smaller then L before rounding. Hence by scaling with $1/\alpha = 10d^2L^{d-2}R$ we get that all coordinates are integers (Lemma 10) and the maximum coordinate has size $10d^2L^{d-1}R$. Plugging in the definition of L gives as upper bound $10d^2R^2$ as asserted. □

By expressing the quantity R in terms of n we can restate Theorem 1 as the following corollary.

Corollary 2 *For a fixed d, every d-dimensional stacked polytope can be realized on an integer grid polynomial in n. The size of the largest z-coordinate is bounded by* $O(n^{3\log(2d)})$*, all other coordinates are bounded by* $O(n^{2\log(2d)})$*.*

Table 1 lists the induced grid bounds for $d = 3, \dots, 10$.

d	exponent largest non-z-coordinate	exponent largest z-coordinate
3	5.17	7.76
4	6	9
5	6.65	9.97
6	7.17	10.76
7	7.62	11.43
8	8	12
9	8.34	12.51
10	8.65	12.97

Table 1: The induced grid bounds in terms of n up to dimension 10.

4 A simple lower bound

In this section we present a simple lower bound. The basis of our construction is the following graph. Take the tetrahedron and stack a vertex in every face, then take the resulting graph and stack again a vertex in every face. We call this graph B_3; see Fig. 4 for an illustration.

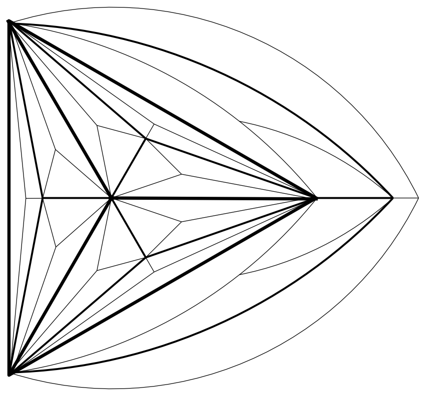

Figure 10: The graph B_3. It has 36 faces and 20 vertices, 12 of them having degree 3.

Lemma 13 *Let P be any embedding of* B_3 *as a stacked 3-polytope. Then there is at least one face with no boundary edge in the orthogonal projection of P into the xy-plane.*

Proof. Let σ be the boundary of $\pi(P)$, which is a polygon. We proceed with a case distinction on the size of σ.
Case 1: σ contains less than 18 vertices.

Every edge on σ is a boundary edge of two faces. Hence in this case we can have at most 34 faces with a boundary edge, but there are 36 faces in B_3. So we have at least two faces without a boundary edge in the projection.

Case 2: σ contains at least 18 vertices.

The cycle σ splits B_3 into two triangulations (interior+cycle and exterior+cycle). Let us now have a look what happens at a degree 3 vertex of B_3 on σ. Two of its incident edges have to be in σ, which means it is the "tip" of an ear in one of the triangulations. Since we have 12 degree 3 vertices and at most two vertices are not on σ, one of the triangulations, say Q, has more than two ears. Now we look at the dual graph of Q in which we removed the vertex for the outer face. Clearly, this graph is connected (since B_3 is 3-connected) and it has as many degree 1 vertices as Q has ears. Hence, there has to be a vertex in the dual graph with degree 3. This means that there is a face in Q, whose adjacent faces are interior faces in Q. Therefore, we have a face in the projection without a boundary edge as asserted. □

There are planar 3-trees that require $\Omega(n^2)$ area for plane straight-line drawings [MNRA11]. Let G_m be such a planar 3-tree with m vertices that needs $\Omega(m^2)$ area. For simplicity we assume that n is a multiple of 36. We glue a copy of $G_{n/36}$ in each of the 36 faces of B_3. This yields another planar 3-tree with n vertices, which we call Γ_n.

Lemma 14 *The embedding of Γ_n as a convex 3-polytope requires a bounding box of $\Omega(n^3)$ volume.*

Proof. Let P be an embedding of Γ_n as a 3-polytope. By restricting the 1-skeleton of P to B_3 and taking the convex hull, we get an embedding of B_3 as a 3-polytope P_B. Due to Lemma 13, there has to be one face, say f, which has no boundary edge on $\pi(P_B)$. The face f defines in $\pi(P)$ a triangle, which contains the graph $G' = G_{n/36}$. The supporting planes of the faces adjacent to f in P_B define a cone C. We denote with C_0 the part $C \setminus P$ that contains the apex of C. Since none of the edges of f are on the boundary of $\pi(P_B)$ we have that $\pi(C_0) = \pi(f)$. Hence, $\pi(P)$ contains a noncrossing drawing of $G_{n/36}$, namely $\pi(G')$, inside $\pi(f)$. Therefore, $\pi(P)$ needs at least area $\Omega(n^2)$.

There was nothing special with choosing the projection into the xy-plane. By the same arguments we can show, that also the projections into the xz- and yz-plane require $\Omega(n^2)$ area.

Let d_x, d_y, d_z denote the dimensions of the bounding box of P along the x-, y-, and z-axis. The volume of the bounding box can be estimated by

$$d_x d_y d_z = \sqrt{d_x^2 d_y^2 d_z^2} = \sqrt{(d_x d_y)(d_x d_z)(d_y d_z)} = \Omega(n^3).$$

□

Acknowledgments.
We thank Alexander Igamberdiev and an anonymous reviewer for helpful comments.

References

[And61] George E. Andrews. A lower bound for the volume of strictly convex bodies with many boundary lattice points. *Trans. Amer. Math. Soc.*, 99:272–277, 1961.

[AŽ95] Dragan M. Acketa and Jovisa D. Žunić. On the maximal number of edges of convex digital polygons included into an $m \times m$-grid. *J. Comb. Theory Ser. A*, 69(2):358–368, 1995.

[BR06] Imre Bárány and Günter Rote. Strictly convex drawings of planar graphs. *Documenta Math.*, 11:369–391, 2006.

[BS10] Kevin Buchin and André Schulz. On the number of spanning trees a planar graph can have. In Mark de Berg and Ulrich Meyer, editors, *Proc. Algorithms – ESA (1)*, volume 6346 of *Lecture Notes in Computer Science*, pages 110–121. Springer, 2010.

[CGT96] Marek Chrobak, Michael T. Goodrich, and Roberto Tamassia. Convex drawings of graphs in two and three dimensions (preliminary version). In *Proc. 12th Symposium on Computational Geometry (SoCG)*, pages 319–328, 1996.

[DG97] Gautam Das and Michael T. Goodrich. On the complexity of optimization problems for 3-dimensional convex polyhedra and decision trees. *Computational Geometry: Theory and Applications*, 8(3):123–137, 1997.

[DLRS10] Jesus De Loera, Jörg Rambau, and Francisco Santos. *Triangulations*, volume 25 of *Algorithms and Computation in Mathematics.* Springer, 2010.

[DS11] Erik D. Demaine and André Schulz. Embedding stacked polytopes on a polynomial-size grid. In *Proc. 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA), San Francsico, 2011*, pages 1177–1187. ACM Press, 2011.

[EG95] Peter Eades and Patrick Garvan. Drawing stressed planar graphs in three dimensions. In Franz-Josef Brandenburg, editor, *Graph Drawing*, volume 1027 of *Lecture Notes in Computer Science*, pages 212–223. Springer, 1995.

[HK92] John E. Hopcroft and Peter J. Kahn. A paradigm for robust geometric algorithms. *Algorithmica*, 7(4):339–380, 1992.

[IS16] Alexander Igamberdiev and André Schulz. A duality transform for constructing small grid embeddings of 3d polytopes. *Comput. Geom.*, 56:19–36, 2016.

[Lov00] László Lovász. Steinitz representations of polyhedra and the Colin de Verdière number. *J. Comb. Theory, Ser. B*, 82:223–236, 2000.

[MNRA11] Debajyoti Mondal, Rahnuma Islam Nishat, Md. Saidur Rahman, and Muhammad Jawaherul Alam. Minimum-area drawings of plane 3-trees. *Journal of Graph Algorithms and Applications*, 15(2):177–204, 2011.

[MRS11] Ares Ribó Mor, Günter Rote, and André Schulz. Small grid embeddings of 3-polytopes. *Discrete & Computational Geometry*, 45(1):65–87, 2011.

[OS94] Shmuel Onn and Bernd Sturmfels. A quantitative Steinitz' theorem. In *Beiträge zur Algebra und Geometrie*, volume 35, pages 125–129, 1994.

[PW13] Igor Pak and Stedman Wilson. A quantitative Steinitz theorem for plane triangulations. `http://arxiv.org/abs/1311.0558`, 2013.

[RG96] Jürgen Richter-Gebert. *Realization Spaces of Polytopes*, volume 1643 of *Lecture Notes in Mathematics.* Springer, 1996.

[RGZ95] Jürgen Richter-Gebert and Günter M. Ziegler. Realization spaces of 4-polytopes are universal. *Bull. Amer. Math. Soc.*, 32:403, 1995.

[Rib06] Ares Ribó Mor. *Realization and Counting Problems for Planar Structures: Trees and Linkages, Polytopes and Polyominoes.* PhD thesis, Freie Universität Berlin, 2006.

[Ryb99] Konstantin A. Rybnikov. Stresses and liftings of cell-complexes. *Discrete & Computational Geometry*, 21(4):481–517, 1999.

[Sch91] Oded Schramm. Existence and uniqueness of packings with specified combinatorics. *Israel J. Math.*, 73:321–341, 1991.

[Sch11] André Schulz. Drawing 3-polytopes with good vertex resolution. *Journal of Graph Algorithms and Applications*, 15(1):33–52, 2011.

[Ste22] Ernst Steinitz. Polyeder und Raumeinteilungen. In *Encyclopädie der mathematischen Wissenschaften*, volume 3-1-2 (Geometrie), chapter 12, pages 1–139. B. G. Teubner, Leipzig, 1922.

[Tar83] Robert Endre Tarjan. Linking and cutting trees. In *Data Structures and Network Algorithms*, chapter 5, pages 59–70. Society for Industrial and Applied Mathematics, 1983.

[Thi91] Torsten Thiele. Extremalprobleme für Punktmengen. Master's thesis, Freie Universität Berlin, 1991.

[Tut60] William T. Tutte. Convex representations of graphs. *Proceedings London Mathematical Society*, 10(38):304–320, 1960.

[Tut63] William T. Tutte. How to draw a graph. *Proceedings London Mathematical Society*, 13(52):743–768, 1963.

[Zic07] Florian Zickfeld. *Geometric and Combinatorial Structures on Graphs.* PhD thesis, Technische Universität Berlin, December 2007.

[Zie95] Günter M. Ziegler. *Lectures on Polytopes.* Springer, 1995.