

MODULI SPACES OF DELZANT POLYTOPES AND SYMPLECTIC TORIC MANIFOLDS

ÁLVARO PELAYO FRANCISCO SANTOS

ABSTRACT. This paper introduces modern geometric combinatorial technology from the theory of triangulations in order to derive results in toric symplectic geometry. In the main part of the paper we prove a number of properties of the space $\mathcal{D}(n)$ of n -dimensional Delzant polytopes. Two highlights are the construction of examples showing that, in contrast with the classical work of Oda in dimension 2, no classification of combinatorially minimal Delzant polytopes can be expected in dimension 3 or higher, and a proof that the space of n -dimensional Delzant polytopes is path-connected. Our proof of the latter is based on the fact that every rational fan can be refined to a unimodular fan, which is a standard technique used for resolution of singularities of toric varieties. In the last part of the paper, using the Delzant correspondence, these results allow us to answer several open questions concerning the moduli space $\mathcal{M}(n)$ of symplectic toric manifolds of dimension $2n$, since this space is isometric to the space of Delzant polytopes. Our results imply that no classification of minimal models of symplectic toric manifolds is plausible in dimension 6 or higher, which answers in the negative a long-standing folklore question originating in Oda's work (1978).

CONTENTS

1. Introduction	2
1.1. Novelty and achievements	2
1.2. From discrete to symplectic geometry and back	2
1.3. Main results	4
1.4. Structure of the paper	5
Acknowledgements	6
2. Delzant polytopes: basic theory and impossibility of an Oda type classification	6
2.1. Polytopes and their normal fans	6
2.2. Stellar subdivisions, desingularization and 2-dimensional classification	11
2.3. Infinitely many minimal Delzant 3-polytopes	15
3. Topology of the space of Delzant polytopes	18
3.1. Two metrics and corresponding topologies on the space of convex bodies	18
3.2. Minkowski paths prove that $\mathcal{D}(n)$ is connected	23
3.3. The stratification of $\mathcal{D}(n)$ by normal equivalence	28
3.4. The fundamental group of $\mathcal{D}(n)$	33
3.5. The CW topology on $\mathcal{D}(n)$.	36
4. Translation to symplectic toric manifolds	38
4.1. Symplectic toric manifolds: basic theory	38
4.2. Geometry and topology of the space of symplectic toric manifolds	41
References	42

2000 *Mathematics Subject Classification*. Primary 52B20, 52C07, 53D20, ; Secondary 53D05, 52B11, 52A20.

Key words and phrases. Delzant polytope, convex polytope, polytope, Delzant correspondence, symplectic toric manifold, moment map, momentum map, moduli space, toric integrable system.

1. INTRODUCTION

1.1. Novelty and achievements. The novelty in this paper is the introduction of modern geometric combinatorial and convex geometric ideas into the study of Delzant polytopes and the moduli spaces they form, which we then incorporate into the framework of symplectic geometry. These techniques include using Hausdorff distance and its close relation to Minkowski sums, focusing on the normal fans of Delzant polytopes rather than on the polytopes themselves, and applying to the normal fans ideas coming from the theory of triangulations of polytopes and vector configurations developed in the past thirty years (see for example the book [12]) but which had not made their way into mainstream symplectic geometry yet.

A *Delzant polytope* in \mathbb{R}^n is a simple full-dimensional polytope with rational edge directions and such that the primitive edge-direction vectors at each vertex form a basis of the lattice \mathbb{Z}^n . Equivalently, it is a polytope with a simplicial and unimodular normal fan (Definition 2.10). We refer to Figure 1 for some examples.




FIGURE 1. Two Delzant polytopes in dimension three. Under the Delzant correspondence, the first one (tetrahedron) corresponds to the toric variety $\mathbb{C}P^3$.

Our two main achievements are:

- We exploit the interplay between fans and polytopes (Figure 2), and use triangulation theory and resolution of singularities of toric varieties in order to construct examples which show that the famous classification of Delzant polygons by Oda [45, Theorem 8.2] cannot have an analogue in dimension three or higher.
- We bring techniques from convex geometry to gain a rather complete understanding of the geometry and topology of the moduli space of n -dimensional Delzant polytopes, extending to arbitrary dimension the results obtained for $n = 2$ by Pelayo, Pires, Ratiu and Sabatini [50, Theorem 2].

While the core ideas of the paper come from geometric combinatorics and convex geometry, they are strongly related to and motivated by developments in symplectic geometry, which build on seminal work of Atiyah, Delzant, Guillemin, Duistermaat, Heckman, Kostant, Sternberg, Weinstein, and others; see [49] for a survey of results in this direction. In the last section of this paper we discuss this connection and the strong implications of the achievements mentioned above in toric symplectic geometry.

1.2. From discrete to symplectic geometry and back. Our combinatorial results have a direct translation into the world of symplectic toric manifolds. As detailed in Definition 4.3 these are quadruples of the form $(M, \omega, \mathbb{T}^n, \mu)$ where \mathbb{T}^n is an n -dimensional




FIGURE 2. The Hirzebruch trapezoid $H_{3/2,1,2}$ with vertices $(0, -\frac{1}{2})$, $(0, \frac{1}{2})$, $(\frac{5}{2}, -\frac{1}{2})$, $(\frac{1}{2}, \frac{1}{2})$ and its corresponding normal fan, with generators $(-1, 0)$, $(0, 1)$, $(0, -1)$, and $(1, 2)$. See Corollary 2.23 for the meaning of its parameters.

torus acting effectively and Hamiltonianly on a $2n$ -dimensional compact connected symplectic manifold (M, ω) with momentum map

$$\mu = (\mu_1, \dots, \mu_n) : M \rightarrow \mathfrak{t}^* \cong \mathbb{R}^n,$$

where \mathfrak{t}^* is the dual Lie algebra of \mathbb{T}^n and there is a choice of isomorphism $\mathfrak{t}^* \cong \mathbb{R}^n$ (see Section 4.1.1 for details and Figure 3 for the simplest example).




FIGURE 3. The only 2-dimensional symplectic toric manifold, up to scaling, $(\mathbb{C}P^1 \cong \mathbb{S}^2, \omega = d\theta \wedge dh, \mathbb{T}^1 = \mathbb{S}^1, \mu(\theta, h) = h)$. The momentum polytope $P = [-1, 1] \subset \mathfrak{t}^* \cong \mathbb{R}$ classifies its \mathbb{S}^1 -equivariant symplectomorphism class. The Delzant correspondence sends the class of $(\mathbb{S}^2, \omega, \mathbb{S}^1, \mu)$ to P .

Often, symplectic toric manifolds are referred to as *toric integrable systems*, since μ is in fact a very special type of finite dimensional integrable Hamiltonian system: one for which all components $\mu_i : M \rightarrow \mathbb{R}$ generate periodic flows of the same period.

The translation of results from the theory of Delzant polytopes to toric symplectic geometry is made via the *Delzant correspondence* (Theorem 4.6), which is a seminal result of Delzant [14] saying that the map

$$\begin{aligned} \mathcal{M}(2n) &\longrightarrow \mathcal{D}(n) \\ [(M, \omega, \mathbb{T}^n, \mu)] &\longmapsto \mu(M) \end{aligned}$$

is a bijection from the set of isomorphism classes of compact connected symplectic toric manifolds of dimension $2n$, which we denote by $\mathcal{M}(2n)$, to the set of Delzant polytopes of dimension n , which we denote by $\mathcal{D}(n)$. Here, two such manifolds are *isomorphic* if there is a toric equivariant symplectomorphism between them preserving the momentum map (Section 4.1.2). The Delzant correspondence establishes a powerful bridge between

the “world of polytopes” and the “world of manifolds”, which is a very rare occurrence in differential geometry. It builds on seminal work of Atiyah [5], Guillemin-Sternberg [27], Horn [32], Kostant [37], and Schur [57].

The use of the Delzant correspondence to work with polytopes instead of manifolds appeared in symplectic geometry around 20 years ago in the papers [47, 48, 42] in a more elementary, from the combinatorial point of view, form. (See also the recent paper [3] for a different approach using convex geometry techniques in symplectic geometry). The definition and moduli space structure on $\mathcal{M}(2n)$ was introduced in [50], and has since then been considered by other authors [22, 21]. It is constructed by pulling back to $\mathcal{M}(2n)$ the symmetric difference distance on $\mathcal{D}(n)$. The metric on $\mathcal{M}(2n)$ is closely related to the famous Duistermaat-Heckman measure [15].

The metrics on $\mathcal{M}(2n)$ and $\mathcal{D}(n)$ obtained in this way have the following nice properties:

- The symmetric difference distance and the Hausdorff distance induce the same topology on $\mathcal{D}(n)$. This allows us to use the latter, which is better-understood in convex geometry, to deduce results about the former, hence about the Duistermaat-Heckman measure.
- The diffeomorphism strata in $\mathcal{M}(2n)$ correspond to the normal equivalence strata in $\mathcal{D}(n)$. Within such strata, continuity with respect to the distance of polytopes is the same as continuity regarding parallel translation of their facets.
- The closures of the strata of $\mathcal{M}(2n)$ correspond to the symplectic toric varieties obtained by exploding lower dimensional invariant submanifolds, which includes blow-ups, but also more complicated explosions.

Having these moduli space structures is useful to study the variation of properties such as continuity of functions defined on these moduli spaces, see Figalli-Pelayo [18] and Figalli-Palmer-Pelayo [19]. We note that other moduli spaces of polygons have been studied by a number of authors, for example see Hausmann-Knutson [29, 30] and Kapovich-Millson [35].

1.3. Main results. Next we describe the main contributions of the paper to geometric combinatorics, convex geometry and symplectic geometry. Sections 2 and 3 do not require any knowledge of symplectic geometry and the results therein are of interest independently of their connection to symplectic geometry, which is worked out in Section 4.

1.3.1. Minimal models of Delzant polytopes and symplectic toric manifolds. In dimension 2 there is a well-known classification of the minimal unimodular complete fans — equivalently, of Delzant polygons— going back to Oda (1978) [45, Theorem 8.2], which we state as Lemma 2.22. This is the main tool used in [50] to study the space $\mathcal{D}(2)$. We show that the failure of such a classification in dimension three and higher is drastic: there are Delzant polytopes with arbitrarily many vertices and facets and which do not admit any blow-down (Theorem 2.26) or, more strongly, which are not refinements of other unimodular fans (Theorem 2.27).

Oda’s Classification implies a classification of 4-dimensional symplectic toric manifolds in terms of finite sequences of blow-ups of three basic minimal models: $\mathbb{C}P^2$, $\mathbb{C}P^1 \times \mathbb{C}P^1$, and the Hirzebruch surfaces. We state the symplectic versions of Theorems 2.26 and Theorem 2.27 as Theorems 4.8 and Theorem 4.10. This answers in the negative the long-standing folklore question in toric geometry of whether an analogue of Oda’s classification of symplectic toric manifolds exists in dimension higher than 4.

1.3.2. Local topology, compactness, and path-connectedness of $\mathcal{D}(n)$ and $\mathcal{M}(2n)$. Indeed, Theorem 3.18 shows that the moduli space $\mathcal{D}(n)$ of n -dimensional Delzant polytopes

is path-connected but not locally compact, and describes its completion under the two natural metrics on it. Theorem 3.19 states the same for the quotient of $\mathcal{D}(n)$ under unimodular equivalence, which we denote by $\widetilde{\mathcal{D}(n)}$, solving in particular a problem posed by Fujita and Ohashi [22, Problem 5.2] in 2018.

Our proof of path-connectedness is algorithmic. It is based on the fact that every pair of unimodular fans admit a common unimodular refinement (Corollary 2.16), which can be computed via the standard resolution of singularities in toric varieties. Yet, contrary to the 2-dimensional case, in dimension three and higher there is no a priori bound on the complexity of the resulting path (Section 3.2.1).

As said above, the Delzant correspondence sends $\mathcal{D}(n)$ bijectively to the space $\mathcal{M}(2n)$ of $2n$ -dimensional symplectic toric manifolds modulo isomorphisms. The quotient $\widetilde{\mathcal{D}(n)}$ by unimodular equivalence then corresponds to quotienting $\mathcal{M}(2n)$ by *weak isomorphisms*; we denote this quotient space by $\widetilde{\mathcal{M}(2n)}$. With this in mind, Theorems 3.18 and 3.19 translate to the symplectic statement Theorem 4.11, which essentially solves [54, Problem 2.42], also posed as [50, Problem 3]. In fact, the path-connectedness part extends the main theorem of [50] from $n = 2$ to arbitrary $n \geq 2$.

1.3.3. *Simple connectedness of $\mathcal{D}(2)$ and $\mathcal{M}(4)$.* We give partial results on the structure of the fundamental group of $\mathcal{D}(n)$, showing that loops with a certain finiteness condition are null-homotopic (Corollary 3.37 and Theorem 3.39). For $n = 2$ all loops have the required property, so $\mathcal{D}(2)$ is simply connected (Corollary 3.45). For $n \geq 3$ loops without the required property exist (Example 3.41), and whether the fundamental group of $\mathcal{D}(n)$ is trivial or not remains an open question (Question 3.46).

The symplectic version of Corollary 3.45 ($\mathcal{M}(4)$ is simply-connected) is Proposition 4.12.

1.3.4. *Stratification of $\mathcal{D}(n)$ and $\mathcal{M}(2n)$.* We study the stratification of $\mathcal{D}(n)$ by normal fans. Let \mathcal{N} be a unimodular fan of dimension n with m rays. We show that the set of Delzant polytopes with normal fan equal to \mathcal{N} form a subspace of $\mathcal{D}(n)$ isometric to a polyhedral relatively open cone of dimension m with a lineality space of dimension n . The faces of the cone correspond to the fans that are refined by \mathcal{N} (Corollary 3.32).

This in particular allows us to give the exact value of the Hausdorff dimension of the space of Delzant n -polytopes with a fixed number of vertices or of facets (Theorem 3.33), thus solving two problems of Fujita and Ohashi [22, Problems 5.3 and 5.5].

1.3.5. *CW topology on $\mathcal{D}(n)$ and $\mathcal{M}(2n)$.* The stratification of the previous paragraph suggests that one can consider $\mathcal{D}(n)$, hence also $\mathcal{M}(2n)$, as embedded into a CW complex. The CW topology obtained in this way is strictly finer than the metric topology (Propositions 3.47 and 3.49) and turns out to be more manageable from a homotopical point of view: with this topology $\mathcal{D}(n)$ and $\mathcal{M}(2n)$ are *weakly contractible* (Corollary 3.51).

1.4. **Structure of the paper.** In order to make the paper accessible for both the combinatorics and symplectic communities we have reviewed the basic concepts of both areas that are needed for the paper. An introduction to polytope theory in general and Delzant polytopes in particular is contained in Sections 2.1 and 2.2 and an introduction to toric symplectic geometry and momentum maps is given in Section 4.1. The rest of Section 2 contains the constructions of 3-dimensional Delzant polytopes that lead to the results related to the impossibility of an Oda type classification, as mentioned in Section 1.3.1 above.

Section 3 deals with the geometry and topology of the space of Delzant polytopes. After introducing some technical tools such as the Hausdorff metric and Minkowski paths in Section 3.1, in Section 3.2 we show that the space of Delzant n -polytopes is path-connected. Section 3.3 studies the stratification of $\mathcal{D}(n)$ by normal fans, and Section 3.4 shows that $\mathcal{D}(2)$ is simply connected. In Section 3.5 we use the aforementioned stratification to endow $\mathcal{D}(n)$ with a CW topology, and show that this new topology is weakly contractible.

After reviewing the basics of toric symplectic geometry in Section 4.1, in Section 4.2 we interpret our combinatorial results in this world.

Acknowledgements. The first author is funded by a BBVA (Bank Bilbao Vizcaya Argentaria) Foundation Grant for Scientific Research Projects with project title *From Integrability to Randomness in Symplectic and Quantum Geometry*. The second author is funded by grant PID2019-106188GB-I00 and PID2022-137283NB-C21 of MCIN/AEI/10.13039/501100011033 and by project CLaPPo (21.SI03.64658) of Universidad de Cantabria and Banco Santander.

We are grateful to the Department of Mathematics, Statistics and Computation at the University of Cantabria for inviting the first author in October 2022 for a visit during which the first ideas of this paper were born. The first author is also thankful to Carlos Beltrán and Fernando Etayo for the hospitality during his visit.

We thank Liat Kessler and Margaret Symington for pointing out to us an error in a previous version of the proof of Theorem 2.27, and Ignasi Mundet for a suggestion leading to the statement of Corollary 3.51, which is more general than the statement we previously had. We also thank an anonymous referee for a very attentive and thorough reading and many useful comments which have helped us improved the paper.

2. DELZANT POLYTOPES: BASIC THEORY AND IMPOSSIBILITY OF AN ODA TYPE CLASSIFICATION

In this section we review some modern tools of geometric combinatorics and convex geometry which we need to prove our main results. We also include several new examples and counterexamples with interesting properties (notably, Theorem 2.26 and Theorem 2.27) which illustrate striking differences between the geometry/topology of the space of Delzant polygons and its higher dimensional analogue. In particular, these examples tell us that no Oda type classification of Delzant n -polytopes is possible in $n \geq 3$.

Throughout the section we emphasize the interplay between fans (more commonly used in algebraic geometry) and polytopes (more commonly used in symplectic geometry), using one or the other depending on how convenient it is for the proofs.

2.1. Polytopes and their normal fans. Because of the close relation which exists between Delzant polytopes and smooth toric varieties, we mainly use the books [8, 17, 23] for background in polytope theory. Occasionally we use [12, 60] for combinatorial aspects, and [24] for convex geometric ones.

2.1.1. Polytopes. Let V be a real vector space of any finite dimension. Let V^* be its dual vector space and let

$$\langle \cdot, \cdot \rangle : V^* \times V \rightarrow \mathbb{R}$$

be the standard pairing.

Definition 2.1. A *polytope* P in V is the closed convex hull of a finite subset $\{p_1, \dots, p_m\} \subset V$, i.e., the smallest convex set containing it or, equivalently,

$$\text{Conv}\{p_1, \dots, p_m\} := \left\{ \sum_{i=1}^m \lambda_i p_i \mid \lambda_i \in [0, 1], \sum_{i=1}^m \lambda_i = 1 \right\}.$$

Many authors call “convex polytope” what we call a polytope. Since we are not concerned with non-convex polytopes we omit the word “convex”, as done for instance in [12, 60].

The *dimension* of a convex set $S \subset V$ is the dimension of its affine span. The set is *full dimensional* if it affinely spans V . (For example, Figure 1 shows two full-dimensional polytopes in $V = \mathbb{R}^3$.)

By definition, a polytope is a compact subset of V . An *extreme point* of a convex subset $C \subset V$ is a point of C which is not in any open line segment contained in C . A polytope is equal to the closed convex hull of its extreme points (Krein-Milman Theorem [38]), and this is the unique minimal description of a polytope as a convex hull of points.¹ Polytopes have also a dual description via inequalities, that is, as finite intersections of half-spaces.

Definition 2.2. A *half-space* in V is any subset which has the form

$$\{p \in V \mid \langle \alpha, p \rangle \leq b\},$$

where $\alpha \in V^* \setminus \{0\}$ and $b \in \mathbb{R}$. We call a *polyhedron* any finite intersection of half-spaces.

By the Minkowski-Weyl Theorem [24, Theorem 14.2] every polytope is a polyhedron and every bounded polyhedron is a polytope.

The language of inequalities allows us to define faces of a polytope, understanding the vector α in each inequality as a normal vector to the corresponding face:

Definition 2.3. Let P be a polytope, or a polyhedron. For each linear functional $\alpha \in V^*$ the *face of P in the direction of α* is

$$F_\alpha(P) := \left\{ p \in P \mid \langle \alpha, p - x \rangle \geq 0 \text{ for } x \in P \right\}.$$

Every face of a polyhedron (respectively, of a polytope) is itself polyhedron (respectively, a polytope), of a certain dimension. The set of all faces is a partially ordered set (*poset*) ranked by dimension. It has a unique maximal element, namely P itself which equals $F_0(P)$. Faces of dimension 0 are called *vertices* of P . Faces of dimension one are called *edges* if they are bounded, which is always the case if P itself is bounded. Faces of dimension $\dim(P) - 1$ are called *facets*.²

A *cone* in V is any subset that is closed under addition and multiplication by nonnegative scalars. A *polyhedral cone* is a cone that is also a polyhedron. Equivalently, it is a finitely generated cone, according to the following formula for the *cone generated by a finite set $v_1, \dots, v_m \in V$* of vectors:

$$\text{Cone}\{v_1, \dots, v_m\} := \left\{ \sum_{i=1}^m a_i v_i \mid a_i \in [0, \infty) \right\}.$$

A polyhedral cone is *simplicial* if it is generated by linearly independent vectors.

¹For a polytope, extreme points are the same as *vertices* (that is, 0-dimensional faces as defined below). For general convex sets, every vertex is an extreme point but the converse is not always true.

²Some authors consider the empty set as a face of P too (of dimension -1), but we do not need that.

The *dual* or *normal cone* of a cone $C \subset V$ is defined as

$$C^\vee := \left\{ \alpha \in V^* \mid \langle \alpha, x \rangle \leq 0 \text{ for } x \in C \right\} \subset V^*.$$

It is easy to show that $(C^\vee)^\vee = C$ for every closed cone, and that the dual of a polyhedral cone is a polyhedral cone.

To every face F of a polytope (or polyhedron) P one naturally associates two closed polyhedral cones, living respectively in V and V^* , and dual to one another: the *tangent cone* and the *normal cone* of a face F of P are the polyhedral cones

$$T(P, F) := \mathbb{R}_{\geq 0}(P - F) = \left\{ \lambda(q - p) \mid q \in P, p \in F, \lambda \geq 0 \right\} \subset V,$$

$$N(P, F) := \left\{ \alpha \in V^* \mid F \subseteq F_\alpha(P) \right\} = \left\{ \alpha \in V^* \mid \langle \alpha, p - q \rangle \geq 0 \text{ for } p \in F, q \in P \right\} \subset V^*.$$

See Figure 4 for an example.




FIGURE 4. Left: a polytope P (a quadrilateral) and a face F of it (a vertex). Right: the tangent and normal cones of F in P .

Definition 2.4. A *fan* in V is a complex of polyhedral cones. That is, a collection \mathcal{N} of polyhedral cones in V that is closed under taking faces and such that if C_1 and C_2 are in \mathcal{N} then $C_1 \cap C_2$ is a common face of both. A fan is called *simplicial* if all its cones are simplicial.

The *normal fan* $\mathcal{N}(P)$ of a polytope P , is the fan consisting of the normal cones of all faces of P . See Figure 2 for an Example.

The cones of $\mathcal{N}(P)$ form again a poset, ranked by dimension, and this poset is anti-isomorphic (that is, isomorphic with the order reversed) to the face poset of P . For example, $\mathcal{N}(P)$ has a unique minimal cone, the normal cone of P itself, which equals $\{0\}$ if and only if P is full-dimensional. When this happens 1-dimensional cones of $\mathcal{N}(P)$ are *rays*, that is, vector half-lines, and they are in bijection to facets of P : each ray in $\mathcal{N}(P)$ is generated by a normal vector to the corresponding facet. In this case P has a unique minimal expression as an intersection of half-spaces, namely the one that uses the half-spaces containing facets of P in their boundary.

The following are some properties that a polyhedron P may or may not have, and the corresponding properties for its normal fan $\mathcal{N}(P)$. We assume P to be full-dimensional and call n the dimension of P and V :

- P is bounded, i.e., a polytope, if and only if its normal fan is *complete*; that is, if every $\alpha \in V^*$ lies in the normal cone of some face.
- P is full-dimensional if and only if its normal fan is *pointed*, that is, if the lineality space of $\mathcal{N}(P)$ is the cone $\{0\}$. Here, the *lineality space* of a cone C is

$$\text{lin}(C) := \left\{ x \in C \mid -x \in C \right\}.$$

It equals the unique maximal linear space contained in C , as well as the unique minimal face of C . In a fan, all cones have the same lineality space. For example, for a polytope (or polyhedron) P the lineality space of $\mathcal{N}(P)$ equals the normal cone of P considered as a face of itself, that is, the space of linear functionals that are constant on P . It is also the linear space orthogonal to the affine span of P .

- A full-dimensional polytope (or polyhedron) P is called *simple* if every vertex lies in exactly n facets. This happens if and only if $\mathcal{N}(P)$ is simplicial.

In dimension two every complete fan is the normal fan of some polygon. In dimension three or higher the same is not true. We say that a complete fan is *polytopal* if it equals the normal fan of some polytope. Otherwise it is *non-polytopal*.

Example 2.5 (A non-polytopal icosahedral fan). Consider the cubeoctahedron, that is, the convex hull of the mid-points of the twelve edges of the cube $[-1, 1]^3 \subset \mathbb{R}^3$. (See, e.g., Example 3.6.17 and Figure 3.47 in [12]). It has six square facets and eight triangular ones. If each square facet is triangulated by inserting one of its two diagonals, we get a triangulated sphere with 20 triangles. A particularly symmetric way of doing this is by choosing the six diagonals to use each of the twelve vertices once, producing a complex isomorphic to the icosahedron. It turns out that the simplicial fan obtained by coning from the origin to each of the 20 triangles is not polytopal, as shown, in a different language, in [12, Example 3.6.17].

Example 2.6 (The smallest non-polytopal fan). Consider the negative standard basis vectors $-e_i$ and the vectors

$$a_i = (e_1 + e_2 + e_3) - e_i, \quad i = 1, 2, 3.$$

Observe that their convex hull is a triangular prism. Triangulate the boundary of the prism inserting diagonals in the quadrilateral faces in a cyclic way: e_1a_2 , e_2a_3 and e_3a_1 . Consider the fan obtained coning from the origin to this triangulation, which has eight cones and six rays. Combinatorially, but not metrically, this fan coincides with the normal fan of the standard 3-cube, that is, the central fan of the regular octahedron. But geometrically this fan is not polytopal; it is closely related to the so-called “mother of all examples” featured prominently in [12]. Every complete fan of dimension n with less than $n - 3$ rays is polytopal ([12, Section 5.5]), so this is the smallest possible non-polytopal fan.

There even exist complete simplicial fans that are not *combinatorially equivalent* to the normal fan of any polytope.³ The smallest one has dimension four, eight rays and 19 maximal cones, and can be derived by coning over Barnette’s sphere [12, Example 9.5.23].

2.1.2. Unimodular fans and Delzant polytopes. From now on we assume that $V = \mathbb{R}^n$. A fan in \mathbb{R}^n is called *rational* if the linear span of every face is a rational subspace. When the fan is pointed (e.g., when it is the normal fan of a full-dimensional polytope) a necessary and sufficient condition for rationality is that every ray contains a non-zero lattice point. For example, the normal fan of a full-dimensional polytope P is rational if, and only if, all facet normals of P have rational directions.

Lemma 2.7. *Let P be a full-dimensional polytope. Then, all edge directions of P are rational if and only if all facet-normal directions of P are rational.*

³In fact, these fans are the ones usually called *non-polytopal* in geometric combinatorics. See, e.g., [12, Section 9.5]. What we call polytopal fans are called *coherent*, *regular*, or *projective* fans, depending on the author.

Proof. The direction of an edge e is the unique orthogonal line to all the normals of facets containing e . The direction normal to a facet F is the unique orthogonal line to all the edges contained in F . \square

Observe that every rational ray R has a unique *primitive generator*, the unique nonzero integer vector in R with relatively prime entries. The same happens in a rational line, except that the primitive generator is now unique only up to sign. We will often abuse language and say that a certain vector $v \neq 0$ “is” a ray in a fan, meaning that it spans a ray. When the ray is rational we will typically assume v to be primitive. The *multiplicity* of a rational simplicial cone C with primitive rays v_1, \dots, v_m is the index of the sublattice generated by $\{v_1, \dots, v_m\}$ as a subgroup of the ambient lattice $\mathbb{R}C \cap \mathbb{Z}^n$. If C is full-dimensional, and hence $m = n$, then the multiplicity of C equals the absolute value of the determinant $\det(v_1, \dots, v_n)$.

Definition 2.8. A *unimodular fan* in \mathbb{R}^n is a rational simplicial full-dimensional fan in which every cone has multiplicity one.

If the cone C_2 is a face of a cone C_1 then the multiplicity of C_2 divides that of C_1 . Hence, in the definition of unimodular only the maximal cones need to be considered.

Proposition 2.9. *For each vertex of a full-dimensional simple polytope (or polyhedron) P , the tangent cone is unimodular if and only if the normal cone is unimodular.*

Proof. Let n be the dimension. Since the tangent cone and the normal cone at a vertex are orthogonal to one another, we can consider the lattice Λ_v generated by the primitive edge vectors at a vertex p . This is a sublattice of \mathbb{Z}^n , of a certain index K , and its dual lattice Λ_v^* is generated by the facet-normals at v , suitably normalized. Hence, Λ_v^* is a superlattice of \mathbb{Z}^n with the same index K . We then have $\Lambda_v = \mathbb{Z}^n$ if and only if $\Lambda_v^* = \mathbb{Z}^n$. That is, the primitive facet normals are a basis of \mathbb{Z}^n if and only if the primitive edge-direction vectors are. \square

The following class of polytopes owe their name to Delzant’s paper [14] (see [25, p. 8]), in which he uncovers a deep and far-reaching relation between symplectic geometry and geometric combinatorics, which we will exploit in the last section of this paper.

Definition 2.10. A *Delzant polytope* of dimension n is a full-dimensional polytope $P \subset \mathbb{R}^n$ whose normal fan is unimodular. Equivalently, it is a full-dimensional polytope satisfying:

- Its normal fan is *rational*. Equivalently, its edge directions are rational;
- It is *simple*: each vertex is contained in exactly n edges or, equivalently, in exactly n edges;
- It is *smooth*: all its cones have multiplicity one. Equivalently, for every vertex v , the primitive vectors along the n edges containing v form a basis of the lattice \mathbb{Z}^n .

Unimodular fans may be non-polytopal, and hence not be the normal fans of any Delzant polytope. In fact, slight modifications of examples 2.5 and 2.6 show this:

- In Example 2.5 all the full-dimensional cones have determinant equal to ± 2 , but their generators lie in the sublattice

$$\left\{ (x, y, z) \in \mathbb{Z}^2 \mid x + y + z \equiv 2 \pmod{2} \right\},$$

of index two. Thus, a linear transformation sending this sublattice to \mathbb{Z}^2 produces a unimodular fan that is still not polytopal. In fact, if this unimodular fan was polytopal it would be the normal fan of a Delzant polytope combinatorially

isomorphic to the dodecahedron, which does not exist: by [13, Theorem 3.1 and Lemma 3.3] every Delzant 3-polytope has at least one triangular or quadrilateral face.

- The fan in Example 2.6 is almost unimodular: it has a unique non-unimodular full-dimensional cone, namely the cone of (a_1, a_2, a_3) which has determinant two. But this cone can be easily divided into three unimodular cones by a stellar subdivision (see definition below) inserting the vector $e_1 + e_2 + e_3 = \frac{1}{2}(a_1, a_2, a_3)$. This gives a unimodular but not polytopal fan, with seven rays and 10 cones.

However, by definition the normal fans of Delzant polytopes are exactly the complete unimodular *and polytopal* fans. For this reason most combinatorial properties of Delzant polytopes can be stated, and are often easier to deal with, via their normal fans.

Remark 2.11. Unimodular fans modulo $\text{GL}(n, \mathbb{Z})$ equivalence are in bijection to smooth toric varieties modulo isomorphism. More generally, rational full-dimensional fans biject to normal toric varieties. In this bijection simplicial fans correspond to varieties that have only finite quotient singularities, complete fans correspond to complete varieties, and polytopal fans to projective ones. Faces of a fan are in bijection with (algebraic) torus orbits, and the multiplicity of each quotient singularity (or orbit of them) equals the multiplicity of the corresponding simplicial cone.

2.2. Stellar subdivisions, desingularization and 2-dimensional classification.

2.2.1. *Stellar subdivision and desingularization.* Let Σ and Σ' be two complete fans in \mathbb{R}^n . We say that Σ' refines Σ if every cone $C' \in \Sigma'$ is contained in some cone $C \in \Sigma$. The fact that cones in a fan intersect properly then implies that every cone of Σ equals a union of cones of Σ' .

A particularly simple and useful method to refine fans is by introducing a single new ray, as follows.

Definition 2.12 (Stellar subdivision, [8, Section 11.1], [17, Definition 2.1]). Let Σ be a fan in \mathbb{R}^n and let $\gamma \in V$ be a vector contained in some cone of Σ . The *stellar subdivision* of Σ by γ is the fan consisting of the cones of Σ not containing γ together with the cones of the form $\text{Cone}\{\gamma \cup C\}$ such that $\gamma \notin C \in \Sigma$ and there is a $C' \in \Sigma$ with $C \cup \{\gamma\} \subset C'$.

Differently put, if we call *star of γ in Σ* the subfan of Σ consisting of cones containing γ , what the stellar subdivision does is to introduce in Σ (the ray generated by) γ by centrally subdividing all cones in the star of γ .

In the case of a simplicial fan, which is our main interest, a more explicit definition of stellar subdivision is possible. Let $C = \text{Cone}\{\alpha_1, \dots, \alpha_k\}$ be the unique minimal cone of Σ containing γ , with $\alpha_1, \dots, \alpha_k$ being ray generators. That is, C is the cone containing γ in its relative interior. Then, every cone containing γ is of the form

$$D = \text{Cone}\{\alpha_1, \dots, \alpha_k, \beta_1, \dots, \beta_l\},$$

for some, perhaps none, generators β_1, \dots, β_l . The stellar subdivision by γ removes each such cone D and inserts instead the k cones

$$D_i := \text{Cone}\{\{\gamma, \alpha_1, \dots, \alpha_k, \beta_1, \dots, \beta_l\} \setminus \{\alpha_i\}\}.$$

The following properties of stellar subdivisions are easy to prove and left to the reader.

Proposition 2.13. *Let Σ be a complete fan in \mathbb{R}^n and let Σ' be obtained from Σ by performing a stellar subdivision at γ .*

- (1) *If Σ is polytopal, then Σ' is polytopal.*

- (2) If Σ is simplicial, then Σ' is simplicial.
- (3) If Σ is complete, then Σ' is complete.
- (4) If Σ is rational and γ is rational, then Σ' is rational.
- (5) If Σ is unimodular and (with the notation above) the α_i are primitive and $\gamma = \sum_{i=1}^k \alpha_i$, then Σ' is unimodular.

The unimodular case is particularly important, so we give it a name:

Definition 2.14. Let Σ be a fan in \mathbb{R}^n and let $C \in \Sigma$ be a unimodular simplicial cone of Σ with primitive generators $\alpha_1, \dots, \alpha_k$. The *blow-up of Σ at C* is the stellar subdivision by the vector $\gamma = \sum_{i=1}^k \alpha_i$. We also say that Σ is obtained from Σ' by a *blow-down*.

The reason why we call this a *blow-up* is that the toric variety of the fan Σ' is the blow-up of the variety of Σ at the torus-invariant orbit corresponding to the cone C . (See [8, Definition 3.3.14 and Proposition 3.3.15], or [17, Definition 7.1]).

Theorem 2.15 ([8, Theorem 11.1.9], [17, Theorem 8.5], [23, Section 2.6]). *Let Σ be any rational complete fan in \mathbb{R}^n . Then there is a sequence of stellar subdivisions producing a unimodular fan Σ' that refines Σ .*

Proof. In a first step we refine Σ to be simplicial. One systematic way of doing this is using the *barycentric subdivision*, which can be expressed as a sequence of stellar subdivisions, one with a γ chosen in the relative interior of each cone of any dimension of Σ , and performed in decreasing order of dimension.

Once we have a simplicial and still rational fan, each cone has a well-defined multiplicity. As mentioned above, the multiplicity of a simplicial cone C with primitive generators $\alpha_1, \dots, \alpha_k$ is the index of $\sum_{i=1}^k \alpha_i \mathbb{Z}$ as a sublattice of the lattice $\text{lin}(C) \cap \mathbb{Z}^n$. This equals the number of lattice points in the half-open cube

$$\sum_{i=1}^k [0, 1) \alpha_i \subset C.$$

Hence, if the multiplicity of C is greater than one, there is a $\gamma \in \mathbb{Z}^n \setminus \{0\}$ which can be expressed as

$$\gamma = \sum_{i=1}^k \lambda_i \alpha_i,$$

with $\lambda_i \in [0, 1)$ for every i . Under these conditions, a stellar subdivision by γ lowers the multiplicity of all the cones involved. More precisely, with the notation after Definition 2.12, each cone D of a certain multiplicity m is substituted by k cones D_i of respective multiplicities $\lambda_i m < m$, for $i \in \{1, \dots, k\}$. Of course, the number of cones increases, but iterating as follows we can make all the resulting cones unimodular: if we call $M > 1$ the maximum multiplicity among the cones of Σ , performing a stellar subdivision with γ chosen in the half-open cube of a cone of multiplicity M we make the number of cones of that multiplicity decrease, and do not introduce cones of higher multiplicity. Hence, eventually we get a fan with maximum multiplicity $M' < M$. Iterating this process, we eventually get a unimodular fan. \square

Observe that by Proposition 2.13, if Σ is polytopal then Σ' is polytopal too. We often use the following immediate consequence of this theorem.

Corollary 2.16. *Let $\Sigma_1, \dots, \Sigma_k$ be rational complete fans in \mathbb{R}^n . Then, there is a unimodular fan Σ' that refines every Σ_i .*

2.2.2. *Corner chopping and blow-ups.* In what follows two subsets X, Y of \mathbb{R}^n are called $\text{AGL}(n, \mathbb{Z})$ -congruent if there exists $c \in \mathbb{R}^n$ and $A \in \text{GL}(n, \mathbb{Z})$ such that $Y = A(X) + c$.

Definition 2.17. The *rational length* of a line segment $I \subset \mathbb{R}^n$ of rational slope is the unique number $\ell \in [0, \infty)$ such that I is $\text{AGL}(n, \mathbb{Z})$ -congruent to an interval of length ℓ on a coordinate axis. It is denoted by $\text{length}(I)$.

Definition 2.18. Let P be a Delzant polytope in \mathbb{R}^n . Let v a vertex of P . Let

$$\left\{ v + tu_i \mid 0 \leq t \leq \ell_i \right\}$$

be the set consisting of edges emanating from v , where the vectors u_1, \dots, u_n generate \mathbb{Z}^n and $\ell_i = \text{length}(u_i)$ for every $i \in \{1, \dots, n\}$. Let $\varepsilon > 0$ be smaller than all the ℓ_i 's. We define the *corner chopping of size ε* of P at v is the polytope P' obtained from P by intersecting it with the half space

$$\left\{ v + \sum_{i=1}^n t_i u_i \mid \sum_{i=1}^n t_i \geq \varepsilon \right\}.$$

Let us translate this operation into the normal fan.

Proposition 2.19. Let P be a Delzant polytope in \mathbb{R}^n . Let v a vertex of P and let P' be obtained from P by a corner chopping of size ε at v . Let $\alpha_1, \dots, \alpha_n$ be the primitive normal vectors of P at the facets containing v and let $a_i = \langle \alpha_i, v \rangle$ be the corresponding right-hand sides for the facet inequalities. Then

$$P' = P \cap \left\{ x \in V \mid \left\langle \sum_{i=1}^n \alpha_i, x \right\rangle \geq \sum_{i=1}^n a_i + \varepsilon \right\}.$$

Proof. With the notation of Definition 2.18, observe that the bases $\{u_i\}_i$ of V and $\{\alpha_i\}_i$ of V^* are dual to one another; that is, $\langle \alpha_i, u_j \rangle = \delta_{i,j}$.

Since any vector $x \in V$ can be written uniquely as $x = v + \sum_{i=1}^n t_i u_i$ we only need to check that the inequality $v + \sum_{i=1}^n t_i u_i \mid \sum_{i=1}^n t_i \geq \varepsilon$ from the definition is equivalent to the inequality $\langle \sum_{i=1}^n \alpha_i, x \rangle \geq \sum_{i=1}^n a_i + \varepsilon$ in the statement. This follows from

$$\left\langle \sum_{i=1}^n \alpha_i, x \right\rangle = \left\langle \sum_{i=1}^n \alpha_i, v + \sum_{j=1}^n t_j u_j \right\rangle = \sum_{i=1}^n \langle \alpha_i, v \rangle + \sum_{i=1}^n t_i = \sum_{i=1}^n a_i + \sum_{i=1}^n t_i. \quad \square$$

That is to say, the normal fan of P' is constructed from that of P by inserting the ray $\sum_{i=1}^n \alpha_i$ in $\mathcal{N}(P, v) = \text{Cone}\{\alpha_1, \dots, \alpha_n\}$.

Corollary 2.20. If a Delzant polytope P' is obtained from another one P by a corner chopping then its normal fan $\mathcal{N}(P')$ is obtained from $\mathcal{N}(P)$ by a blow-up at a full-dimensional cone $C \in \mathcal{N}(P)$.

The converse is also true, in the following sense:

Corollary 2.21. Let P' be a Delzant polytope in \mathbb{R}^n . If its normal fan $\mathcal{N}(P')$ is the blow-up of a certain unimodular fan Σ at a full-dimensional cone then there is a Delzant polytope P with $\mathcal{N}(P) = \Sigma$ and of which P' is a corner chopping.

Proof. Let $\alpha_1, \dots, \alpha_n$ and $\alpha_0 = \sum_{i=1}^n \alpha_i$ be the rays of Σ' involved in the blow-up, and let b_0, b_1, \dots, b_n be the corresponding right-hand sides in the inequality description of P' . Observe that all the cones C_i in the definition of blow-up automatically have the same determinant as the cone C generated by $\alpha_1, \dots, \alpha_n$. In particular the blow-down fan Σ is unimodular, since Σ' is.

The facet of P' corresponding to α_0 is an $(n - 1)$ -simplex, since it has n vertices, corresponding to the n cones C_i in the definition of blow-up. Removing from the inequality description of P' the inequality $\langle \alpha_0, x \rangle \geq b_0$ gives a polytope P in which that facet disappears and the adjacent n facets (corresponding to the rays $\alpha_1, \dots, \alpha_n$) are enlarged, producing a unique new vertex v , the intersection of the hyperplanes $\{\langle \alpha_i, x \rangle = b_i\}$, $i \in \{1, \dots, n\}$. The normal cone of v , generated by $\{\alpha_1, \dots, \alpha_n\}$, is unimodular as said above. Hence, P is Delzant and P' is obtained from it as in Proposition 2.19, that is, by a corner chopping. \square

2.2.3. Classification of Delzant polygons. A two-dimensional fan is uniquely determined by its set of rays, since each two-dimensional cone is bounded by two rays and the rays must appear in their cyclic order. This easily implies that all two-dimensional fans are polytopal: if $\alpha_1, \dots, \alpha_n$ are generators for the rays of the fan Σ and we identify \mathbb{R}^2 and $(\mathbb{R}^2)^*$ via the standard scalar product, then the polytope defined by the linear inequalities $\langle \alpha_i, x \rangle \leq \|\alpha_i\|$ has Σ as its normal fan.

Any unimodular two-dimensional fan can be obtained by the following recursive recipe.

Lemma 2.22 (Oda [45, Theorem 8.2]. See also [8, Theorem 10.4.3],[17, Theorem 6.6],[23, Section 2.5 and Notes to Chapter 2]). *Let Σ be a complete unimodular fan in \mathbb{R}^2 .*

- (1) *If Σ has only three rays then it is $\text{GL}(2, \mathbb{Z})$ -equivalent to the fan of the projective plane, with rays $(1, 0)$, $(0, 1)$ and $(-1, -1)$.*
- (2) *If Σ has four rays then it is $\text{GL}(2, \mathbb{Z})$ -equivalent to the fan of the Hirzebruch surface \mathbb{H}_k , with rays $(0, -1)$, $(0, 1)$, $(-1, 0)$, and $(1, k)$, for some $k \in \mathbb{Z}_{\geq 0}$.*
- (3) *If Σ has more than four rays then it is an iterated blow-up of some unimodular fan with four rays.*

The following translation of Lemma 2.22 into Delzant polygons is formulated in [36, Lemma 2.16].

Corollary 2.23. *Let P be a Delzant polygon in \mathbb{R}^2 .*

- (1) *If P has exactly three edges then there exists a unique $\lambda > 0$ such that P is $\text{AGL}(2, \mathbb{Z})$ -congruent to the Delzant triangle*

$$P_\lambda := \{ (x_1, x_2) \in \mathbb{R}^2 \mid x_1 \geq 0, x_2 \geq 0, x_1 + x_2 \leq \lambda \}.$$

- (2) *If P has exactly with 4 edges, then there exists $k \in \mathbb{Z}_{\geq 0}$ and positive numbers $a, b > 0$ with $2a > bk$ such that P is $\text{AGL}(2, \mathbb{Z})$ -congruent to a Delzant polygon obtained from the Hirzebruch trapezoid*

$$H_{a,b,k} := \left\{ (x_1, x_2) \in \mathbb{R}^2 \mid -\frac{b}{2} \leq x_2 \leq \frac{b}{2}, 0 \leq x_1 \leq a - kx_2 \right\},$$

by a sequence of s corner choppings (as defined in Definition 2.18).

- (3) *If P has $4 + s$ edges for $s \geq 1$ then it is obtained from a Hirzebruch trapezoid by a sequence of s corner choppings.*

The parameters in $H_{a,b,k}$ have the following interpretation: b is the height of the trapezoid, a is its average width, and the integer k governs the difference in slopes between the two non-parallel edges (with the trapezoid being a rectangle if $k = 0$). Observe also that $H_{a,b,1}$ is a corner chopping of (a translation of) the triangle $P_{a+\frac{b}{2}}$. See Figure 2 for a picture of $H_{3/2,1,2}$.

Remark 2.24 (The maximum area of a Delzant polygon). One interesting question is how big can the area of a Delzant polygon be in terms of its number k of vertices

(equivalently, edges) and its perimeter ℓ . Here, the perimeter is the sum of lengths of edges, with length as in Definition 2.17. The following results answer this question asymptotically.⁴ Then:

- (1) Every Delzant polygon with k vertices and perimeter ℓ has area bounded above by $\phi^{2k}\ell^2$, where $\phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio [7, Theorem 28].
- (2) There are Delzant polygons with k vertices and perimeter ℓ with area $\geq \phi^{2k/3}\ell^2$ [7, Example 30].

A lower bound in terms of k and ℓ alone makes no sense: one can construct Delzant polygons with a pair of parallel edges and arbitrarily many additional edges. Making the two parallel edges very long and the rest of edges very short one can keep ℓ and k arbitrarily large, yet make area go to zero.

Remark 2.25. Via Delzant’s correspondence, Lemma 2.22 and Corollary 2.23 imply a classification of (compact) toric integrable systems in dimension 4. We refer to [51, 52] for an analogue classification in the *semitoric* case, and [33, Theorem 1.3] and [34, Theorem 1.1] for a discussion of these concepts (fan, etc.) in that case. The topology on the moduli space of semitoric integrable systems was defined by Palmer [46].

2.3. Infinitely many minimal Delzant 3-polytopes. In dimension three and higher, corner chopping alone can certainly not produce all Delzant polytopes from a finite list of combinatorial types of them. Consider, for example, the prism obtained as the Cartesian product $P \times I$ of a Delzant m -gon P and a segment I . This is a Delzant 3-polytope with $m + 2$ facets, none of them triangular (unless $m = 3$), hence it is not a corner chopping of any Delzant polytope with less facets.

However, in this example we can still say that if $m \geq 5$ then P is a corner chopping of a Delzant $(m - 1)$ -gon P' , and then $P \times I$ is an “edge chopping” of $P' \times I$. Equivalently, we have that the normal fan of $P \times I$ is a blow-up of the fan of $P' \times I$ at a non-fulldimensional cone, in the sense of Definition 2.14.

The following two constructions show more drastically that one should not hope any analogue of Lemma 2.22 or Corollary 2.23 in dimension three or higher. In the first example we show not only that there are arbitrarily large 3-dimensional unimodular fans that do not admit blow-downs, but also that such examples cannot be considered “sporadic”, since any given fan can be refined to one of them. In the following proof we call *valency* of a ray in a 3-dimensional fan the number of full-dimensional cones containing it.

Theorem 2.26. *Let Σ be any rational complete fan in \mathbb{R}^3 . Then, there is a unimodular fan Σ' that refines Σ and such that Σ' is not the blow up of any unimodular fan. If Σ is polytopal then Σ' can be obtained polytopal too.*

Proof. By Theorem 2.15 we may assume without loss of generality that Σ is already unimodular. We proceed in two steps, illustrated in Figure 5.

Let Σ'' be the unimodular fan obtained from Σ by blowing up all its 3-dimensional cones. These blow-ups double the valency of each ray of Σ . Since all rays in a complete 3-dimensional fan have valency at least three, all the original rays of Σ have valency at least six in Σ'' , while the rays introduced by the blow-ups have valency three.

⁴The statements in [7] assume the Delzant polytopes in question have only integer vertices, but the proofs go through without this assumption.

We now define Σ' subdividing each full-dimensional cone $\text{Cone}\{\alpha_1, \alpha_2, \alpha_3\} \in \Sigma''$ into the following six cones:

$$\begin{array}{ccc} \text{Cone}\{\alpha_1, \alpha_{12}, \alpha_{13}\}, & \text{Cone}\{\alpha_2, \alpha_{12}, \alpha_{23}\}, & \text{Cone}\{\alpha_3, \alpha_{13}, \alpha_{23}\}, \\ \text{Cone}\{\alpha_{123}, \alpha_{12}, \alpha_{13}\}, & \text{Cone}\{\alpha_{123}, \alpha_{12}, \alpha_{23}\}, & \text{Cone}\{\alpha_{123}, \alpha_{13}, \alpha_{23}\}, \end{array}$$

where $\alpha_{ij} := \alpha_i + \alpha_j$ and $\alpha_{123} := \alpha_1 + \alpha_2 + \alpha_3$. Since the α_{ij} and α_{123} are primitive and the six 3-dimensional cones are unimodular, Σ' is again a unimodular fan. The rays of Σ'' keep their valency in Σ' , while the new rays have valency three if they are of the form α_{123} and valency eight if they are of the form α_{ij} . (For the latter, let α_{ij} be one of them; Since Σ' is complete, α_{ij} lies in the common boundary of two three-dimensional cones of Σ' , and is contained in four of the six three-dimensional cones refining each).

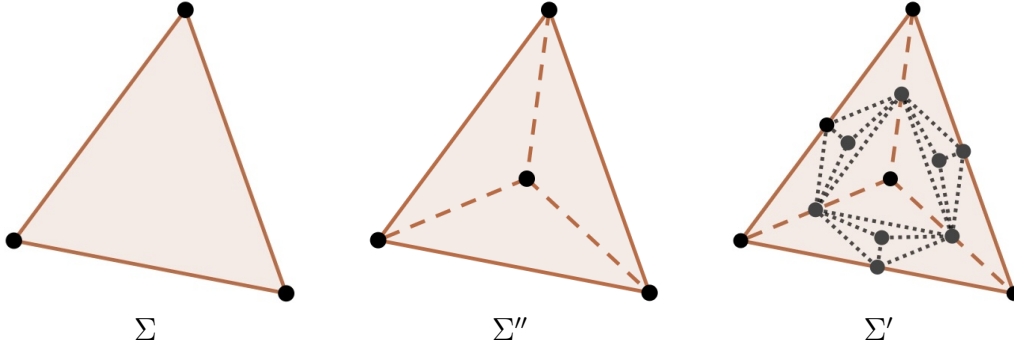


FIGURE 5. The refinement constructed in the proof of Theorem 2.26. Left: A 3-dimensional cone of the original fan Σ , represented as a triangle, via an intersection with an affine plane. Center: its blow-up in Σ'' , consisting of three 3-dimensional cones separated via three new two-dimensional cones (dashed lines) with one new ray (center dot). Right: In Σ' each 3-cone from Σ'' is subdivided into six 3-cones, introducing several new 2-cones (pointed lines) and new rays of two types: one along each 2-cone of Σ'' , denoted α_{ijk} in the proof, and one in each 3-cone of Σ'' , denoted α_{ij} in the proof.

To show that Σ' is not the blow-up of any unimodular fan we observe that the new ray in a blow-up of a three-dimensional complete fan has valency three or four, depending on whether the blown-up cone is 3-dimensional or 2-dimensional. That is, only rays of valency three or four in Σ' are candidates to be blown-down. By the discussion above, Σ' has no rays of valency four, and two types of rays of valency three:

- (1) The rays introduced in the blow-ups from Σ to Σ'' . Let α be one of them, and let β, γ, δ be its neighbors in Σ'' , so that $\alpha = \beta + \gamma + \delta$. The three neighbors of α in Σ' are $\alpha + \beta$, $\alpha + \gamma$, and $\alpha + \delta$. Since their sum is 4α , α cannot be blown down. (Combinatorially it may look like it can, but the cone $\text{Cone}\{\alpha + \beta, \alpha + \gamma, \alpha + \delta\}$ obtained by this supposed blow-down has multiplicity four, not one).
- (2) The rays of the form α_{ijk} introduced when going from Σ'' to Σ' . These have the same problem; blowing down α_{123} would produce the cone $\text{Cone}\{\alpha_{12}, \alpha_{12}, \alpha_{23}\}$, of multiplicity 2. \square

In dimension two, a unimodular fan is a blow-up if and only if it is the refinement of another unimodular fan. In dimension three, the fans of the previous theorem do not admit blow-downs but they still are refinements of unimodular fans. Thus, one could perhaps hope that an analogue of Lemma 2.22 can exist in dimension three except the refinements needed in the classification are a more complicated list than only blow-ups. The following examples kill this hope, in any dimension $n \geq 3$:

Theorem 2.27. *In any dimension $n \geq 3$ there are complete unimodular polytopal fans with arbitrarily many rays which are not proper refinements of any other unimodular fan.*

Proof. We first construct them for $n = 3$. Let \mathcal{V}_k be the following set of $k + 6$ rays in \mathbb{R}^3 , for an arbitrary $k \in \mathbb{N}$:

$$\begin{aligned}\alpha_+ &= (0, 0, 1), & \alpha_- &= (k + 1, 1, -1), \\ \beta_{-2} &= (0, -1, 0), & \beta_{-3} &= (1, 0, 0), \\ \beta_i &= (i, 1, 0), & i &\in \{-1, \dots, k\}.\end{aligned}$$

We consider the following complete simplicial fan \mathcal{N}_k in it, with $2k + 8$ maximal cones:

$$\begin{aligned}\text{Cone}\{\beta_{-3}, \alpha_+, \alpha_-\}, & \quad \text{Cone}\{\beta_k, \alpha_+, \alpha_-\}, \\ \text{Cone}\{\beta_{i-1}, \beta_i, \alpha_+\}, & \quad \text{Cone}\{\beta_{i-1}, \beta_i, \alpha_-\}, \quad i = -2, \dots, k.\end{aligned}$$

Observe that the configuration and the fan are invariant under the unimodular involution

$$\begin{aligned}\mathbb{R}^3 &\longrightarrow \mathbb{R}^3 \\ (x_1, x_2, x_3) &\longmapsto (x_1 + (k + 1)x_3, x_2 + x_3, -x_3)\end{aligned}$$

which exchanges α_+ and α_- and fixes the rest of \mathcal{V}_k .




FIGURE 6. The configuration \mathcal{V}_k and complete fan \mathcal{N}_k proving Theorem 2.27, pictured for $k = 5$. The picture shows it sliced with the hyperplane $H = \{x_3 = 0\}$ containing all generators except α_+ and α_- ; the vector $\alpha_+ + \alpha_- \in H$ is not in \mathcal{V}_k , but it is shown to represent the intersection of H with the 2-face $\text{Cone}\{\alpha_+, \alpha_-\} \in \mathcal{N}_k$.

One crucial property of this construction is that all the β_i lie in the plane $H := \{x_3 = 0\}$ and the other two generators α_+ and α_- lie in opposite sides of it. This allows us to picture \mathcal{V}_k and \mathcal{N}_k in two dimensions, by intersecting them with H , as we do in Figure 6. In fact, the fan \mathcal{N}_k can be understood as follows: start with the unimodular two-dimensional complete fan of Figure 6 (including the ray $\alpha_+ + \alpha_-$), which is clearly unimodular. Take the cones of it with both α_+ and α_- , which gives a unimodular complete fan in \mathbb{R}^3 with four maximal cones surrounding the ray $\alpha_+ + \alpha_-$. Then, blow-down $\alpha_+ + \alpha_-$; that is, remove the four maximal cones containing it and insert $\text{Cone}\{\beta_{-3}, \alpha_+, \alpha_-\}$ and $\text{Cone}\{\beta_k, \alpha_+, \alpha_-\}$ instead. This description shows that \mathcal{N}_k is indeed a complete unimodular fan.

We claim that if \mathcal{N}' is a complete unimodular fan refined by \mathcal{N}_k then $\mathcal{N}' = \mathcal{N}_k$. For this:

- (1) The rays α_+ and α_- must be used in \mathcal{N}' since each of them is the only element of \mathcal{V}_k in a certain open half-plane (namely, the half-planes $x_3 > 0$ and $x_3 < 0$).

- (2) This implies that the cone $C = \text{Cone}\{\alpha_+, \alpha_-\}$ must be in \mathcal{N}' too. Indeed, C must be contained in some cone C' of \mathcal{N}' and since the generators of C are used in \mathcal{N}' and C' is simplicial, C must be a face of C' .
- (3) The only β_i 's that form a unimodular cone together with $\text{Cone}\{\alpha_+, \alpha_-\}$ are β_{-3} and β_k . In order to check this, observe that the determinant of $\text{Cone}\{\beta_i, \alpha_+, \alpha_-\}$ as a 3-dimensional cone is the same as that of $\text{Cone}\{\beta_i, \alpha_+ + \alpha_-\}$ as a 2-dimensional one. We conclude that

$$\text{Cone}\{\beta_{-3}, \alpha_+, \alpha_-\}, \text{Cone}\{\beta_k, \alpha_+, \alpha_-\} \in \mathcal{N}'.$$

- (4) Once we know this, all the other maximal cones must be of the form $\text{Cone}\{\beta_i, \beta_j, \alpha_\pm\}$ for some i, j and some choice of sign for α_\pm . Hence, the rest of the fan must consist of two-dimensional cones in H joined to both α_+ and α_- , and the cones used in H must be unimodular themselves. An easy inspection of Figure 6 shows that the only way to complete the unimodular fan is exactly as \mathcal{N}_k does.

Once the statement is proved for $n = 3$ we prove it for any $n \geq 4$ by induction. Hence, we assume that for every k there is a full-dimensional complete polytopal unimodular fan Σ_k with at least k rays in \mathbb{R}^{n-1} that is not the proper refinement of any unimodular fan. We consider the *suspension* of Σ_k defined as:

$$\text{susp}(\Sigma_k) := \left\{ \sigma \times \{0\} \mid \sigma \in \Sigma_k \right\} \cup \left\{ \sigma \times [0, \infty) \mid \sigma \in \Sigma_k \right\} \cup \left\{ \sigma \times (-\infty, 0] \mid \sigma \in \Sigma_k \right\}.$$

Put differently, from each cone $\text{Cone}\{\alpha_1, \dots, \alpha_\ell\}$ in Σ_k (including the 0 cone, generated by the empty set) we have the following three cones in $\text{susp}(\Sigma_k)$:

$$\text{Cone}\{\alpha'_1, \dots, \alpha'_\ell\}, \quad \text{Cone}\{\alpha'_1, \dots, \alpha'_\ell, e_n\}, \quad \text{Cone}\{\alpha'_1, \dots, \alpha'_\ell, -e_n\},$$

where $\alpha'_i := \alpha_i \times \{0\}$, $i = 1, \dots, \ell$.

$\text{susp}(\Sigma_k)$ is an n dimensional complete unimodular fan with 2 more rays than Σ_k . It is not a refinement of any fan since any such fan would need to be of the form $\text{susp}(\Sigma')$ where Σ_k refines Σ' . It is polytopal since it is the normal fan of $P \times [0, 1]$ for any polytope P with normal fan Σ_k . \square

Remark 2.28. The example in the first part of the proof of Theorem 2.27 is based on the *one-point suspension* construction, described for example in [12, Section 4.2.5]. We say that a configuration \mathcal{V} of dimension n is a one-point suspension if there are two vectors $\alpha_+, \alpha_- \in \mathcal{V}$ such that $\mathcal{V}' := \mathcal{V} \setminus \{\alpha_+, \alpha_-\} \cup \{\alpha_+ + \alpha_-\}$ is $(n - 1)$ -dimensional. Every simplicial fan \mathcal{N} with generators contained in \mathcal{V} (called a “triangulation of \mathcal{V} ” in the language of [12]) is obtained as the example in the proof: starting with a triangulation \mathcal{N}' of \mathcal{V}' do the following to each maximal cone C of it:

- If $\alpha_+ + \alpha_-$ is a generator of C , substitute $\alpha_+ + \alpha_-$ by α_+ and α_- as generators.
- If $\alpha_+ + \alpha_-$ is not a generator of C , cone C separately to α_+ and to α_- .

The suspension used in the second part of the proof is the case $\alpha_+ + \alpha_- = 0$.

3. TOPOLOGY OF THE SPACE OF DELZANT POLYTOPES

In this section we study natural geometric structures (metric, topology) on the spaces of Delzant polytopes of any dimension.

3.1. Two metrics and corresponding topologies on the space of convex bodies.

Following Gruber [24] we denote by $\mathcal{C}(n)$ the space of all convex bodies (i.e., compact convex sets) in \mathbb{R}^n and by $\mathcal{C}_p(n)$ the subset of *proper* ones; that is, those with non-empty interior or, equivalently, with positive volume. Also, we denote by $\mathcal{P}(n)$ the subspace of

perhaps not full-dimensional polytopes, and by $\mathcal{P}_{\mathbb{Q}}(n)$ the subspace of polytopes with rational edge directions (or, equivalently by Lemma 2.7, those with rational facet normals).

Definition 3.1. We consider the following two functions

$$\delta^V, \delta^H : \mathcal{C}(n) \times \mathcal{C}(n) \longrightarrow \mathbb{R}_{\geq 0},$$

called respectively the *symmetric difference distance* and the *Hausdorff distance*:

$$\begin{aligned} \delta^V(P, Q) &:= \text{Vol}(P \setminus Q) + \text{Vol}(Q \setminus P), \\ \delta^H(P, Q) &:= \max \left\{ \max_{x \in P} \text{dist}(x, Q), \max_{y \in Q} \text{dist}(y, P) \right\}, \end{aligned}$$

where

$$\text{dist}(x, Q) = \min_{y \in Q} \|x - y\|.$$

Observe that for δ^V to be a distance function we need to restrict ourselves to \mathcal{C}_p since in it every two improper convex bodies are at distance zero. Alternatively, we can consider δ^V defined on $\mathcal{C}_p \cup \{0\}$, where “0” denotes the class of all non-proper convex bodies.

The following is immediate.

Proposition 3.2. *The pairs $(\mathcal{C}_p \cup \{0\}, \delta^V)$ and (\mathcal{C}, δ^H) are metric spaces.*

Restricted to \mathcal{C}_p , δ^H and δ^V are not equivalent as distances; for example, the sequence of rectangles $[0, m] \times [0, 1/m^2]$, with $m \in \mathbb{N}$ is a Cauchy sequence with respect to δ^V since the volume of the rectangles, hence of their symmetric differences, goes to zero. However, it is not Cauchy with respect to δ^H since the Hausdorff distance between any two of these rectangles is at least 1.

Pelayo-Pires-Ratiu-Sabatini [50, Section 2.1] consider the space of Delzant polytopes with the topology induced by δ^V . We here will primarily work with δ^H since it relates better to normal fans. Fortunately, the two distances define the same topology, which gives us a well-defined space of Delzant polytopes:

Proposition 3.3 (δ^H versus δ^V). *The following statements hold.*

- (1) *Restricted to $\mathcal{C}_p(n)$ the Hausdorff distance and the symmetric difference induce the same topology, which is locally compact.*
- (2) *$\mathcal{C}_p(n)$ is not complete for any of the two distances. Its completion for δ^H is $\mathcal{C}(n)$ and its completion for δ^V is $\mathcal{C}_p(n) \cup \{0\}$.*

Proof. That δ^H and δ^V induce the same topology is proven in [58] (see also [50, Remark 10]), and that the topology induced by δ^H is locally compact is [24, Theorem 6.4].

That the completion of δ^H is $\mathcal{C}(n)$ appears in [24, Theorem 6.4], where it is moreover proved that $\mathcal{C}(n)$ is boundedly compact (stronger than being complete).

That the completion for δ^V is $\mathcal{C}_p(n) \cup \{0\}$ can be derived as follows (the case of $n = 2$ is [50, Theorem 7], with a different proof): let $(P_m)_{m \in \mathbb{N}}$ be a sequence of polytopes that is Cauchy in δ^V . Since we only need to find a convergent subsequence, there is no loss of generality in assuming that for every $m \in \mathbb{N}$ and $m' > m$

$$\delta^V(P_m, P_{m'}) \leq 2^{-m}.$$

For each m let

$$Q_m = \bigcup_{m' \geq m} P_{m'}.$$

This is not a convex body but it is a measurable set and this is enough in order for δ^V to be well-defined. We have that

$$\delta^V(Q_m, P_m) \leq \sum_{m' > m} \delta^V(P_{m'}, P_m) \leq 2^{-m+1}.$$

Hence, the sequence $(P_m)_{m \in \mathbb{N}}$ converges if, and only if, the decreasing sequence $(Q_m)_{m \in \mathbb{N}}$ converges. We claim that the latter converges to

$$Q = \overline{\bigcap_{m \in \mathbb{N}} Q_m},$$

and that Q either has volume zero or is a convex body. For the first part:

$$\begin{aligned} \delta^V(Q_m, Q) &= \text{Vol}(Q_m \setminus Q) \\ &\leq \sum_{m' \geq m} \text{Vol}(Q_{m'} \setminus Q_{m'+1}) \\ &\leq \sum_{m' \geq m} \text{Vol}(P_{m'} \setminus P_{m'+1}) \\ &\leq 2^{-m+1}. \end{aligned}$$

For the second part, suppose that Q has positive volume, so its interior is not empty. We first claim that the interior of Q is convex. For this, let p and q be two interior points, so that there are balls $B(p, \varepsilon)$ and $B(q, \varepsilon)$ contained in Q , that is, contained in infinitely many of the original P_m . Since the original sequence was Cauchy, $B(p, \varepsilon/2)$ and $B(q, \varepsilon/2)$ must be contained in all but perhaps finitely many of the P_m and, by convexity of P_m , the same happens for the segment pq .

We now claim that Q equals the closure of its interior, and hence it is convex. For this, let $q \in Q$ and $p \in Q^\circ$. Choosing as before a ball $B(p, \varepsilon)$ contained in Q , we have that q lies in the closure of the interior of $\text{Conv}\{B(p, \varepsilon) \cup \{q\}\}$, which in turn is contained in the closure of the interior of Q .

Finally, since Q is a closed convex set with non-empty interior and of finite volume, it must be bounded, hence a convex body. \square

In what follows we denote by $\widetilde{\mathcal{C}}(n), \widetilde{\mathcal{C}}_p(n)$ respectively the quotients of $\mathcal{C}(n), \mathcal{C}_p(n)$ by the action of $\text{AGL}(n, \mathbb{Z})$, with the quotient topology, and consider the map

$$(1) \quad \begin{aligned} \widetilde{\delta^V} : \widetilde{\mathcal{C}}_p(n) \times \widetilde{\mathcal{C}}_p(n) &\longrightarrow \mathbb{R}_{\geq 0} \\ ([P], [Q]) &\longmapsto \inf_{\substack{P' \in [P], \\ Q' \in [Q]}} \delta^V(P', Q'), \end{aligned}$$

where $[P]$ and $[Q]$ denote the orbits of P and Q .

Lemma 3.4. *For any proper convex body $P \in \mathcal{C}_p(n)$ and any bounded set K there are only a finite amount of elements $\sigma \in \text{AGL}(n, \mathbb{Z})$ such that $\sigma(P) \subset K$.*

Proof. Since P is proper, there is an $N \in \mathbb{N}$ such that P contains an affine basis B of $\frac{1}{N}\mathbb{Z}^n$. Now, $\frac{1}{N}\mathbb{Z}^n$ is invariant under $\text{AGL}(n, \mathbb{Z})$ and knowing $\sigma|_B$ completely determines each $\sigma \in \text{AGL}(n, \mathbb{Z})$, so the number of possible σ with $\sigma(P) \subset K$ is bounded above by the number of $(n+1)$ -tuples of points of $\frac{1}{N}\mathbb{Z}^n$ in K , which is finite. \square

The following example, which as far as we know is new, shows that Lemma 3.4 does not hold for improper convex bodies. In particular, it shows that the quotient topology on $\widetilde{\mathcal{C}}(n)$ is not Hausdorff.

Proposition 3.5. *Any neighborhood of the origin contains infinitely many distinct $\text{AGL}(2, \mathbb{Z})$ images of the segment S joining $(0, 0)$ and $(\phi, -1)$.*

Proof. Consider the unimodular matrices

$$\begin{pmatrix} F_k & F_{k+1} \\ F_{k+1} & F_{k+2} \end{pmatrix}$$

where the F_k are the Fibonacci numbers. They take S to segments of lengths going to zero as k goes to infinity, so that there is an infinite amount of them in every neighborhood of the origin. \square

The following result shows that the infimum in expression (1) is really a minimum and it implies that the distance between different orbits is positive. This is also proved in [22, Proposition 4.6] but we want to emphasize that our proof is considerably shorter.

Lemma 3.6. *Let $P, Q \in \mathcal{C}_p(n)$ be proper convex bodies. Then $\widetilde{\delta^V}([P], [Q])$ equals $\delta^V(P, Q')$ for some $Q' \in [Q]$.*

Proof. We first prove the following claim: let $m, M \in \mathbb{R}_{\geq 0}$, with $M > m$, and consider the boxes $B_m = [-m, m]^n$ and $B_M = [-M, M]^n$. Then, every convex body K not contained in B_M has at least a fraction of

$$(M - m)^n / (M + m)^n$$

of its volume outside B_m .

To prove the claim, let $p = (p_1, \dots, p_n) \in K \setminus B_M$, so that at least one coordinate of p is strictly larger than M in absolute value. Suppose without loss of generality that $p_1 > M$. Let

$$K_0 = K \cap \left\{ (x_1, \dots, x_n) \in \mathbb{R}^n \mid x_1 = m \right\},$$

and let K_1 be the following bounded affine cone with apex at p :

$$K_1 := \left\{ p + \lambda(q - p) \mid q \in K_0, \lambda \geq 0 \right\} \cap \left\{ (x_1, \dots, x_n) \in \mathbb{R}^n \mid -m \leq x_1 \right\}.$$

Observe that the fraction of volume of K_1 contained in $\{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_1 \geq m\}$ is exactly $(M - m)^n / (M + m)^n$. The claim follows then from the fact that

$$\begin{aligned} K \cap B_m &\subset K_1 \cap \left\{ (x_1, \dots, x_n) \in \mathbb{R}^n \mid x_1 \leq m \right\}, \\ K \cap (B_M \setminus B_m) &\supset K_1 \cap \left\{ (x_1, \dots, x_n) \in \mathbb{R}^n \mid x_1 \geq m \right\}. \end{aligned}$$

To prove the lemma first observe that

$$\delta^V(P', Q') = \text{Vol}(P') + \text{Vol}(Q') - 2 \text{Vol}(P' \cap Q').$$

Since volume is invariant under $\text{AGL}(n, \mathbb{Z})$, the first two terms are independent of the representatives P' and Q' chosen. Hence, we only need to understand

$$\sup_{P' \in [P], Q' \in [Q]} \text{Vol}(P' \cap Q').$$

Invariance under the action of $\text{AGL}(n, \mathbb{Z})$ also shows that this supremum equals

$$\sup_{Q' \in [Q]} \text{Vol}(P \cap Q').$$

If $\text{Vol}(P, Q') = 0$ for every $Q' \in [Q]$ then there is nothing to prove, so we assume that we have a representative $Q_0 \in [Q]$ with $\text{Vol}(P \cap Q_0) > 0$. Let $m, M \in \mathbb{R}$ be such that P is contained in the box $o + B_m$ for some point o and such that

$$\frac{(M - m)^n}{(M + m)^n} \geq \frac{\text{Vol}(Q) - \text{Vol}(P \cap Q_0)}{\text{Vol}(Q)}.$$

Such an m exists since P is bounded, and M exists since the left-hand side in this equation goes to 1 as M goes to infinity with m fixed, while the right-hand-side is smaller than 1. Observe that this right-hand side equals the fraction of volume of Q_0 outside P .

By the claim, any convex body not contained in $o + B_M$ has less fraction of its volume intersecting $o + B_m$ (and hence even less intersecting P) than Q_0 . Now, every Q' in $[Q]$ has the volume of Q and only finitely many of them are contained in B_M by Lemma 3.4, so only finitely many polytopes in $[Q]$ can have a larger intersection with P than Q_0 . \square

Corollary 3.7 ([22, Theorem 3.2]). $(\widetilde{\mathcal{C}}_p(n), \delta^V)$ is a metric space.

Proof. Lemma 3.6 implies that different orbits $[P]$ and $[Q]$ have strictly positive distance $\widetilde{\delta}^V([P], [Q])$, and it also implies the triangular inequality:

$$\begin{aligned} \widetilde{\delta}^V([P], [Q]) + \widetilde{\delta}^V([Q], [R]) &= \inf_{P' \in [P]} \delta^V(P', Q) + \inf_{R' \in [R]} \delta^V(Q, R') \\ &= \inf_{\substack{P' \in [P], \\ R' \in [R]}} (\delta^V(P', Q) + \delta^V(Q, R')) \\ &\geq \inf_{\substack{P' \in [P], \\ R' \in [R]}} (\delta^V(P', R')) \\ &= \widetilde{\delta}^V([P], [R]). \end{aligned}$$

\square

Remark 3.8. There are several obstacles to generalizing these results on the quotient metric to the Hausdorff distance δ^H . The first one is that one would like (at least if we want to understand the completion of $\widetilde{\mathcal{C}}(n)$ as we do next) to deal also with improper convex bodies but, as shown by Proposition 3.5, the quotient space $\widetilde{\mathcal{C}}(n)$ is not Hausdorff.

But even for proper convex bodies, where the quotient topology of the metric spaces $(\mathcal{C}(n), \delta^V)$ and $(\mathcal{C}(n), \delta^H)$ coincides, the distance $\widetilde{\delta}^H$ between two different orbits can be zero. Consider for example the standard triangle $P = \text{Conv}\{(0, 0), (1, 0), (0, 1)\}$ and the unit square $Q = [0, 1]^2$. Then

$$P' = \text{Conv}\{(0, 0), (1, 0), (k, 1)\} \quad \text{and} \quad Q' = \text{Conv}\{(0, 0), (1, 0), (k-1, 1), (k, 1)\}$$

are representatives with $\delta^H(P', Q') = 1/\sqrt{k^2 + 1}$, which goes to zero as k goes to infinity.

Related to this, Macbeath proved [40, Theorem 1] that the quotient of $\mathcal{C}_p(n)$ by the action of $\text{AGL}(n, \mathbb{R})$ is compact and metrizable, with the distance between P and Q defined essentially as $|\log(\text{Vol}(Q'')/\text{Vol}(Q'))|$ where Q' and Q'' are the biggest and smallest (respectively) elements of the orbit of Q with $Q' \subset P \subset Q''$.

The following is an analogue of Proposition 3.3 for quotient spaces by the $\text{AGL}(n, \mathbb{Z})$ -action. Part (1) is [22, Theorem 3.3], and part (2) allows us to solve Problem 5.2 in the same paper (Theorem 3.19).

Proposition 3.9. *The following statements hold.*

- (1) The metric topology of $(\widetilde{\mathcal{C}_p(n)}, \widetilde{\delta^V})$ coincides with the quotient topology induced on $\widetilde{\mathcal{C}_p(n)}$ by the action of $\text{AGL}(n, \mathbb{Z})$ on $\mathcal{C}_p(n)$.
- (2) The metric space $(\widetilde{\mathcal{C}_p(n)}, \widetilde{\delta^V})$ is not complete. Its completion is $\widetilde{\mathcal{C}_p(n)} \cup \{0\}$.

Proof. In order to show that the topologies coincide we need to show that the identity map between the quotient and metric topologies in $\widetilde{\mathcal{C}_p(n)}$ is continuous and open. For this, observe that a ball $B([P], \varepsilon) \subset \widetilde{\mathcal{C}_p(n)}$ with respect to the metric $\widetilde{\delta^V}$ lifts, by Lemma 3.6, to $\cup_{P' \in [P]} B(P', \varepsilon)$; since the latter is open in $\mathcal{C}_p(n)$ the former is open in the quotient topology. Conversely, if a $\widetilde{U} \subset \widetilde{\mathcal{C}_p(n)}$ is open in the quotient topology then \widetilde{U} lifts to an open and $\text{AGL}(\mathbb{Z}, n)$ -invariant set U in $\mathcal{C}_p(n)$. For each P with $[P] \in \widetilde{U}$ we have that there is a ball $B(P, \varepsilon)$ contained in U and, by invariance of U , the same ε satisfies that $B(P', \varepsilon) \subset U$ for every $P' \in [P]$. Hence, \widetilde{U} contains the ball $B([P], \varepsilon)$.

For the completion, we extend $\widetilde{\delta^V}$ to $\widetilde{\mathcal{C}_p(n)} \cup \{0\}$ defining

$$\widetilde{\delta^V}([P], 0) = \text{Vol}(P).$$

Since we know $\mathcal{C}_p(n) \cup \{0\}$ to be the completion of $(\mathcal{C}_p(n), \delta^V)$ (Proposition 3.3, (2)) we only need to show that every Cauchy sequence in $\widetilde{\mathcal{C}_p(n)}$ lifts to a Cauchy sequence in $\mathcal{C}_p(n)$.

Let $([P_k])_{k \in \mathbb{N}}$ be a Cauchy sequence and, by taking a subsequence if needed, assume that

$$\widetilde{\delta^V}([P_k], [P_{k+1}]) \leq 2^{-k}, \quad \forall k \in \mathbb{N}.$$

Thanks to Lemma 3.6, we can then choose one by one representatives P'_k in each orbit $[P_k]$ in such a way that for every k we have

$$\delta^V(P'_k, P'_{k+1}) = \widetilde{\delta^V}([P_k], [P_{k+1}]) \leq 2^{-k}, \quad \forall k \in \mathbb{N},$$

which gives us a Cauchy sequence $(P'_k)_{k \in \mathbb{N}}$ in $\mathcal{C}_p(n)$. □

3.2. Minkowski paths prove that $\mathcal{D}(n)$ is connected. We now formally introduce the spaces of Delzant polytopes.

Definition 3.10.

- (i) We denote by $\mathcal{D}(n)$ the space of all Delzant n -polytopes in \mathbb{R}^n , with the topology induced by either δ^H or δ^V .
- (ii) We denote by $\widetilde{\mathcal{D}(n)}$ the quotient of $\mathcal{D}(n)$ by the action of $\text{AGL}(n, \mathbb{Z})$, with the quotient topology.

In order to analyze the topology of $\mathcal{D}(n)$ and its relation to $\mathcal{C}(n)$ the *Minkowski sum* operation is very useful. Recall that for given sets $P, Q \subset \mathbb{R}^n$ their Minkowski sum equals

$$P + Q := \left\{ x + y \mid x \in P, y \in Q \right\}.$$

Minkowski sum preserves convexity, boundedness, and compactness so it is well-defined on $\mathcal{C}(n)$. Together with Minkowski sum we consider multiplication by a scalar and, more generally, *Minkowski combinations*. That is, for given $\lambda, \mu \in [0, \infty)$ and $P, Q \in \mathcal{C}(n)$ we consider

$$\lambda P + \mu Q := \left\{ \lambda x + \mu y \mid x \in P, y \in Q \right\}.$$

Minkowski sum is very well behaved with respect to *support functions*. Recall that the support function $h_C : (\mathbb{R}^n)^* \rightarrow \mathbb{R}$ of a convex body $C \subset \mathbb{R}^n$ is defined by

$$h_C(\alpha) := \max_{x \in C} \langle \alpha, x \rangle.$$

Proposition 3.11 ([24, Theorem 6.1, Proposition 6.2, Theorem 4.3]).

- (1) *With the operations of Minkowski addition and multiplication by a scalar $\mathcal{C}(n)$ is a semigroup (with identity $\{0\}$) and a cone.*
- (2) $h_{\lambda P + \mu Q} = \lambda h_P + \mu h_Q$.
- (3) *A function $h : (\mathbb{R}^n)^* \rightarrow \mathbb{R}$ is the support function of some (and then a unique) convex body if and only if it is positive homogeneous and subadditive. That is, if and only if*

$$\forall \alpha, \beta \in (\mathbb{R}^n)^* \text{ and } \forall \lambda, \mu \geq 0 : \quad h(\lambda \alpha) = \lambda h(\alpha), \quad h(\lambda \alpha + \mu \beta) \leq \lambda h(\alpha) + \mu h(\beta).$$

That is to say, algebraically, Minkowski combinations allow us to think of $\mathcal{C}(n)$ as the subcone of positive homogeneous and subadditive functions in the vector space of functions from \mathbb{R}^n to \mathbb{R} . This analogy carries over to the topology of $\mathcal{C}(n)$ under the following setting. Let us call *uniform norm*, and denote $\|f\|_\infty$, of a positively homogeneous continuous function $f : (\mathbb{R}^n)^* \rightarrow \mathbb{R}$ the supremum norm restricted to the unit sphere. That is:

$$\|f\|_\infty := \sup_{\|\alpha\|_2=1} |f(\alpha)|.$$

Then we have that:

Theorem 3.12 ([56, Theorem 1.8.11]). *For any two convex bodies $C, D \in \mathcal{C}(n)$ with support functions $h_C, h_D : (\mathbb{R}^n)^* \rightarrow \mathbb{R}$ we have*

$$\delta^H(C, D) = \|h_C - h_D\|_\infty = \max_{\|\alpha\|_2=1} \left(\max_{x \in C} \langle \alpha, x \rangle - \max_{x \in D} \langle \alpha, x \rangle \right).$$

Corollary 3.13. *Let $C(\mathbb{S}^{n-1})$ be the space of continuous functions from the unit sphere $\mathbb{S}^{n-1} \subset \mathbb{R}^n$ to \mathbb{R} , with the supremum norm. Then, $\mathcal{C}(n)$ is homeomorphic (in fact, isometric) to the subspace of $C(\mathbb{S}^{n-1})$ consisting of functions whose positive homogeneous extension is subadditive.*

The support function of a polytope is linear within each normal cone and, in fact, this characterizes polytopes: a convex body C is a polytope if and only if its support function h_C is piecewise linear. Together with the fact that $h_{P+Q} = h_P + h_Q$, this implies that when P and Q are polytopes $P + Q$ is also a polytope and the normal fan of $\mathcal{N}(P + Q)$ is the common refinement of $\mathcal{N}(P)$ and $\mathcal{N}(Q)$ (that is, cones in $\mathcal{N}(P + Q)$ are the intersections of a cone from $\mathcal{N}(P)$ and one from $\mathcal{N}(Q)$).

The main property of Minkowski sums that we need is the following:

Lemma 3.14. *For every $P, Q \in \mathcal{C}(n)$ the map*

$$\begin{aligned} [0, \infty)^2 &\longrightarrow \mathcal{C}(n) \\ (\lambda, \mu) &\longmapsto \lambda P + \mu Q, \end{aligned}$$

is Lipschitz with respect to δ^H , hence uniformly continuous.

Proof. Let K be the radius of a ball centered at the origin and containing both P and Q and consider two Minkowski sums $A = \lambda P + \mu Q$ and $B = \lambda' P + \mu' Q$, with $\lambda, \lambda', \mu, \mu' \in [0, \infty)$.

Assume that the Hausdorff distance $\delta^H(A, B)$ is attained by the distance to B from a certain point $\lambda x + \mu y \in A$, where $x \in P$ and $y \in Q$. Since the point $\lambda' x + \mu' y$ is in B we

have that

$$\begin{aligned}
\delta^H(A, B) &= \delta^H(\lambda P + \mu Q, \lambda' P + \mu' Q) \\
&= \text{dist}(\lambda x + \mu y, \lambda' x + \mu' y) \\
&\leq \|\lambda x + \mu y - \lambda' x - \mu' y\| \\
&\leq |\lambda - \lambda'| \|x\| + |\mu - \mu'| \|y\| \\
&\leq K(|\lambda - \lambda'| + |\mu - \mu'|) \\
&= K\|(\lambda, \mu) - (\lambda', \mu')\|_1,
\end{aligned}$$

where the last equality is by definition of the norm $\|\cdot\|_1$, which is equivalent to the standard norm in $[0, \infty)^2$. \square

Throughout this paper a *path* in a topological space X is a continuous map $f : [a, b] \rightarrow X$ from a compact interval $[a, b] \subset \mathbb{R}$, with $a < b$. The points $f(a), f(b) \in X$ are the extreme points of the path and we often say that f is a *path from $f(a)$ to $f(b)$* .

Corollary 3.15. *For every $P, Q \in \mathcal{C}(n)$ the map*

$$\begin{aligned}
[0, 1] &\longrightarrow \mathcal{C}(n) \\
\lambda &\longmapsto \lambda P + (1 - \lambda)Q
\end{aligned}$$

is a path from P to Q in $\mathcal{C}(n)$.

Definition 3.16. We call the path of Corollary 3.15 the *Minkowski segment* from P to Q and denote it by PQ .

The Corollary implies that $\mathcal{C}(n)$ is path-connected and, in fact, contractible. (For the latter, fix a convex body K and consider the homotopy $(P, t) \mapsto (1-t)P + tK$). The same holds for any subset of $\mathcal{C}(n)$ closed under Minkowski addition and scaling, for example for $\mathcal{C}_p(n)$, $\mathcal{P}(n)$, and $\mathcal{P}_{\mathbb{Q}}(n)$.

It does not directly imply that $\mathcal{D}(n)$ is path-connected, since the Minkowski sum of two Delzant polytopes is in general not Delzant, or even simple. E.g., the Minkowski sum of the standard triangle with vertices $(0, 0)$, $(1, 0)$, $(0, 1)$ and its reflection on the vertical axis is a non-Delzant pentagon. However, there is an easy way around this issue.

Lemma 3.17. *Let P and Q be two polytopes in \mathbb{R}^n and suppose that the normal fan of P refines that of Q then*

$$\mathcal{N}(\lambda P + (1 - \lambda)Q) = \begin{cases} \mathcal{N}(P) & \text{if } \lambda \neq 0, \\ \mathcal{N}(Q) & \text{if } \lambda = 0. \end{cases}$$

In particular, if P and Q are both in $\mathcal{D}(n)$ then the Minkowski segment between them stays in $\mathcal{D}(n)$.

Proof. $\mathcal{N}(\lambda P + (1 - \lambda)Q)$ is the common refinement of $\mathcal{N}(\lambda P)$ and $\mathcal{N}((1 - \lambda)Q)$. \square

Part (2) of our next Theorem 3.18 and part (1) of Theorem 3.19 were proved in [50, Theorem 2] for the case of $n = 2$ using the classification Corollary 2.23. We here extend both results to arbitrary n and to $\widehat{\mathcal{D}}(n)$.

Theorem 3.18.

- (1) $\mathcal{D}(n)$ is path-connected.
- (2) $\mathcal{D}(n)$ is not locally compact, but it is dense in $\mathcal{C}_p(n)$. Hence, its completion is the same as that of $\mathcal{C}_p(n)$ (namely $\mathcal{C}_p(n) \cup \{0\}$ for δ^V and $\mathcal{C}(n)$ for δ^H).

Proof. In order to prove part (1) we simply note that for given $P_1, P_2 \in \mathcal{D}(n)$ there is, by Corollary 2.16, a Delzant polytope Q whose normal fan refines that of $P_1 + P_2$, hence it refines the normal fans of both P_1 and P_2 . By Lemma 3.17, concatenating the Minkowski segments from P_1 to Q and from Q to P_2 we get a path from P_1 to P_2 that stays in $\mathcal{D}(n)$.

In part (2), the fact that $\mathcal{D}(n)$ is neither complete nor locally compact is proved in [50, Theorem 2] for $n = 2$ and the proof carries over for higher n .

In order to show that $\mathcal{D}(n)$ is dense in $\mathcal{C}_p(n)$ let $K \in \mathcal{C}_p(n)$ and let us show how to approach K via Delzant polytopes, in two steps:

- K can be approached via (perhaps not Delzant) rational polytopes. This is classical and easy to prove. For example, let

$$K_m := \text{Conv} \left(K \cap \frac{1}{m} \mathbb{Z}^n \right),$$

which is a rational polytope contained in K for each $m \in \mathbb{N}$. Since every neighborhood of each $x \in K$ contains points of K_m for every sufficiently large m , we have that

$$\lim_{m \rightarrow \infty} \delta^H(K_m, K) = 0.$$

Here we are implicitly using compactness of K , which allows us to cover K by finitely many arbitrarily small neighborhoods of points $x \in K$.

- Every rational polytope Q can be approached by Delzant polytopes. Let P be a Delzant polytope such that $\mathcal{N}(P)$ refines $\mathcal{N}(Q)$, which exists by Theorem 2.15. Then the path $\lambda P + (1 - \lambda)Q$ from P to Q consists only of Delzant polytopes with the normal fan of P , except at its origin $\lambda = 0$ where it equals Q . \square

Theorem 3.19.

- (1) The metric topology of $(\widetilde{\mathcal{D}(n)}, \widetilde{\delta^V})$ coincides with the quotient topology of $\mathcal{D}(n)$ by the action of $\text{AGL}(n, \mathbb{Z})$.
- (2) $\widetilde{\mathcal{D}(n)}$ is path-connected.
- (3) $(\widetilde{\mathcal{D}(n)}, \widetilde{\delta^V})$ is not locally compact, but it is dense in $\widetilde{\mathcal{C}_p(n)}$. Hence, its completion is the same as that of $\widetilde{\mathcal{C}_p(n)}$, namely $\widetilde{\mathcal{C}_p(n)} \cup \{0\}$.

Proof. Part (1) follows from the same result for $\mathcal{C}(n)$ and $\widetilde{\mathcal{C}(n)}$ (Proposition 3.9) and part (2) is obvious from path-connectedness of $\mathcal{D}(n)$.

Part (3) is also immediate from the same result for $\mathcal{D}(n)$ together with Lemma 3.6. \square

The last part of Theorem 3.19 solves a problem posed by Fujita-Ohashi [22, Problem 5.2] in 2018. We later solve Problems 5.3 and 5.5 of the same paper.

3.2.1. *Local and global complexity: comments on the proof of Theorem 3.18.* Let us call a path in $\mathcal{C}(n)$ a *Minkowski path* if it is a finite concatenation of Minkowski segments.

Our proof of connectedness for $\mathcal{D}(n)$ shows that any two Delzant polytopes P_1 and P_2 can be connected by a Minkowski path $P_1 Q P_2$ consisting of two segments. Metrically, this path can be made as close to the Minkowski segment from P_1 to P_2 as one wishes, via the following trick. As in the proof of the theorem, let Q be a Delzant polytope whose normal fan refines those of P_1 and P_2 . Now, for each $\varepsilon \in (0, 1)$ let

$$Q_\varepsilon = (1 - \varepsilon)P_1 + \varepsilon Q,$$

which has the same normal fan as Q but is very close to P_1 if ε is very small. Then, the path $P_1 Q_\varepsilon P_2$, which is contained in $\mathcal{D}(n)$, is very close to the Minkowski segment $P_1 P_2$.

However, our proof does not give a bound for the following two measures of complexity.

Local complexity. Let us call *complexity of the Minkowski segment* PQ the number of rays of $\mathcal{N}(P)$ that are not in $\mathcal{N}(Q)$ plus the number of rays of $\mathcal{N}(Q)$ that are not in $\mathcal{N}(P)$. Equivalently, this is the number of rays of the intermediate normal fan (which is the same in the interior of the path) that are not common to P and Q . Let us call *local complexity of a Minkowski path* the maximum complexity of its Minkowski segments.

For example, a Minkowski segment between two Delzant polytopes has complexity 1 if and only if P is a blow-up of Q , or viceversa. Hence, a Minkowski path has local complexity 1 if and only if it is a sequence of blow-ups and blow-downs. Our proof of Theorem 3.18 does not give any bound on the local complexity of the path, but the following result does:

Theorem 3.20 (Weak Factorization Theorem [59, Theorem A], see also [4]). *Every two unimodular fans can be joined by a sequence of blow-downs and blow-ups.*

Corollary 3.21. *Every two Delzant polytopes can be joined by a Minkowski path of local complexity one.*

Observe that our proof of path-connectedness is elementary and self-contained, while the Weak Factorization Theorem and its proof are far from elementary, as shown by its history of claimed proofs that turned out to be incorrect.

Remark 3.22. There is also the Strong Factorization statement, asserting that every two unimodular fans can be joined by a sequence of blow-ups followed by a sequence of blow-downs. This was conjectured by Hironaka (1960) [31] and Oda (1978) [45]. Morelli (1996) [43] claimed to have a proof, but gaps in the proof were later discovered. The recent preprint [2] contains a proof of a version where stellar subdivisions (Definition 2.12, sometimes called weighted blow-ups) and their inverses are allowed. That is, the paper proves that any two rational complete simplicial fans admit a common stellar subdivision refinement.

The weighted versions of the Factorization Theorems are important in the (toric) *Minimal Model Program*. In this program one starts with a quasiprojective toric variety X and a given divisor D and wants to modify (X, D) to a toric variety \tilde{X} and a *nef* divisor \tilde{D} via “elementary moves”. At the level of fans, the elementary moves allowed are combinatorial and correspond to *bistellar flips*, where a bistellar flip is the composition of a stellar coarsening (i.e., a weighted blow-down) and a stellar refinement with the same center γ . See [8, Chapter 15] or [20].

In dimension 2, Corollary 3.21 easily follows from Lemma 2.22, as follows:

- (1) By the Lemma, we can go from any Delzant polygon other than a triangle to a quadrilateral by a sequence of blow-downs. Such quadrilaterals are $\mathrm{GL}(\mathbb{Z}, 2)$ -equivalent to Hirzebruch quadrilaterals $H_{a,b,k}$.
- (2) The Hirzebruch quadrilaterals with consecutive parameters k and $k + 1$ can be joined by a blow-up followed by a blow-down (see [50, Figure 4]).
- (3) The Hirzebruch quadrilateral with $k = 1$ is a blow-up of the standard Delzant triangle

$$\left\{ (x_1, x_2) \mid x_1 \geq 0, x_2 \geq 0, x_1 + x_2 \leq 1 \right\}.$$

- (4) Any two different Delzant triangles can be joined by a sequence of blow-ups and blow-downs. This follows easily from the fact that $\mathrm{GL}(2, \mathbb{Z})$ is generated by the transposition of coordinates plus the matrix $\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$.

In fact, these statements (without the last one) are essentially the proof of path-connectedness of $\widetilde{\mathcal{D}(n)}$ in [50, Theorem 5].

Global complexity. By *global complexity of a path* we mean the total number of rays used along the whole path apart of those at the end-points. The global complexity of the Minkowski path from any Delzant polygon to the quadrilateral that it refines is zero, so only the path between two quadrilaterals needs to be considered.

Let P and Q be the two Delzant quadrilaterals that we want to join by a path, and consider their normal fan $P + Q$. In dimension two there is a unique minimal way to desingularize $\mathcal{N}(P + Q)$, namely introducing the rays that belong to the *Hilbert basis* of each cone. The size of the Hilbert basis of a two-dimensional rational cone is bounded by its determinant, so we can get a path from P to Q with global complexity bounded by the sum of the determinants of the cones of $P + Q$. This can be arbitrarily large if P and Q are very “skew” with respect to one another, but we can consider this skewness to be the natural parameter if we want to measure the “global distance” between $\mathcal{N}(P)$ and $\mathcal{N}(Q)$.

3.3. The stratification of $\mathcal{D}(n)$ by normal equivalence.

3.3.1. *The secondary fan of a vector configuration, and the secondary cone of a fan \mathcal{N} .* We say that two polytopes P and Q are *normally equivalent* if they have the same normal fan. Normal equivalence gives a natural stratification of the space of all polytopes. Restricted to $\mathcal{P}_{\mathbb{Q}}(n)$ and hence also to $\mathcal{D}(n)$, it has countably many strata.

To make this more clear, let $\mathcal{V} = \{\alpha_1, \dots, \alpha_m\} \subset (\mathbb{R}^n)^*$ be a finite number of vectors *positively spanning* $(\mathbb{R}^n)^*$, that is, with $\text{Cone } \mathcal{V} = (\mathbb{R}^n)^*$. We will later assume that the α_i are integer and primitive, but for the time being we don't need that.

Let us denote by $\mathcal{P}(\mathcal{V}) \subset \mathcal{P}(n)$ the space of all polytopes $\{P(\mathcal{V}, b) \mid b \in \mathbb{R}^{\mathcal{V}}\}$ where

$$(2) \quad P(\mathcal{V}, b) := \left\{ x \in \mathbb{R}^n \mid \langle \alpha_i, x \rangle \leq b_i, i \in \{1, \dots, m\} \right\},$$

and $\mathbb{R}^{\mathcal{V}} \cong \mathbb{R}^m$ denotes the space of possible right-hand sides of the defining inequalities. $\mathcal{P}(\mathcal{V})$ contains the full-dimensional polytopes whose facet normals are contained in \mathcal{V} , but it also contains lower-dimensional polytopes and, in fact, it contains all points. Indeed, for each $p \in \mathbb{R}^n$, taking $b_i = \langle \alpha_i, p \rangle$ we have that $P(\mathcal{V}, b) = \{p\}$.

Not every vector $b \in \mathbb{R}^{\mathcal{V}}$ defines a polytope, since the corresponding system of inequalities may be infeasible. Also, if one of the equations in the system defining $P(\mathcal{V}, b)$ happens to be redundant then increasing the value of the corresponding b_i does not change the polytope. Hence, we can assume without loss of generality that the b we use equals the support function of $P(\mathcal{V}, b)$ at the vectors α_i of \mathcal{V} . With this assumption we have a bijection between $\mathcal{P}(\mathcal{V})$ and the subset of $\mathbb{R}^{\mathcal{V}}$ obtained by restricting to \mathcal{V} the support functions of polytopes in $\mathcal{P}(\mathcal{V})$. Since we show below (Theorem 3.31) that this bijection is in fact an isometry, we can slightly abuse notation and identify $\mathcal{P}(\mathcal{V})$ with that subset. This subset is a cone, whose explicit description can be derived from part (3) of Proposition 3.11 with the help of Farkas' Lemma.

Lemma 3.23 (Farkas' Lemma [60, Proposition 1.7]). *Let $\alpha_1, \dots, \alpha_m \in (\mathbb{R}^n)^*$ and $b \in \mathbb{R}^m$. Then, the system of inequalities $\langle \alpha_i, x \rangle \leq b_i, i \in \{1, \dots, m\}$, is infeasible if and only if there are $\lambda_1, \dots, \lambda_m \geq 0$ such that*

$$\sum_{i=1}^m \lambda_i \alpha_i = 0, \quad \sum_{i=1}^m \lambda_i b_i < 0.$$

Theorem 3.24. Let $\mathcal{V} = \{\alpha_1, \dots, \alpha_m\} \subset (\mathbb{R}^n)^*$ be a finite number of vectors positively spanning $(\mathbb{R}^n)^*$. The set $\mathcal{P}(\mathcal{V}) \subset \mathbb{R}^{\mathcal{V}}$ of vectors b that are support functions of polytopes with normals contained in \mathcal{V} is the cone defined by the following two types of inequalities:

- (1) For each nonnegative linear dependence $\sum_{i=1}^m \lambda_i \alpha_i = 0$ among the vectors in \mathcal{V} , (that is, linear dependence with all $\lambda_i \geq 0$) the inequality $\sum_{i=1}^m \lambda_i b_i \geq 0$.
- (2) For each equality $\sum_{i=1}^m \lambda_i \alpha_i = \alpha_j$ with all $\lambda_i \geq 0$, the inequality $\sum_{i=1}^m \lambda_i b_i \geq b_j$.

Moreover, only the linear dependences and combinations with minimal support (which are finitely many) are needed.

Proof. For each $b \in \mathbb{R}^{\mathcal{V}}$, let us denote

$$P(b) := \bigcap_{i=1}^m \{x \in \mathbb{R}^n \mid \langle \alpha_i, x \rangle \leq b_i, i \in \{1, \dots, m\}\}.$$

That the inequalities are satisfied for every $b \in \mathcal{P}(\mathcal{V})$ follows from part (3) of Proposition 3.11. Indeed, if we have $\sum_{i=1}^m \lambda_i \alpha_i = 0$ for some nonnegative λ_i 's then:

$$0 = h_{P(b)}(0) = h_{P(b)}\left(\sum_{i=1}^m \lambda_i \alpha_i\right) \leq \sum_{i=1}^m \lambda_i h_{P(b)}(\alpha_i) = \sum_{i=1}^m \lambda_i b_i.$$

Similarly, if $\sum_{i=1}^m \lambda_i \alpha_i = \alpha_j$ then

$$b_j = h_{P(b)}(\alpha_j) = h_{P(b)}\left(\sum_{i=1}^m \lambda_i \alpha_i\right) \leq \sum_{i=1}^m \lambda_i h_{P(b)}(\alpha_i) = \sum_{i=1}^m \lambda_i b_i.$$

Conversely, suppose that we have a certain $b \notin \mathcal{P}(\mathcal{V})$. That is, either the polytope $P(b)$ is empty, or it is not empty but there is a j for which b_j is strictly larger than $h_{P(b)}(\alpha_j)$.

If the first thing happens, that is, if the system $\bigcap_{i=1}^m \{x \in \mathbb{R}^n \mid \langle \alpha_i, x \rangle \leq b_i\}$ is infeasible, then Farkas' Lemma (Lemma 3.23) directly says that there is a linear dependence $\sum_{i=1}^m \lambda_i \alpha_i = 0$ for which the inequality $\sum_{i=1}^m \lambda_i b_i \geq 0$ is violated. If the second happens, then the following system is infeasible:

$$\langle \alpha_i, x \rangle \leq b_i \quad \forall i \in \{1, \dots, m\}, \quad \langle -\alpha_j, x \rangle \leq -b_j.$$

Hence, Farkas' Lemma tells us that there are nonnegative $\lambda_1, \dots, \lambda_m, \mu$ such that $\sum_{i=1}^m \lambda_i \alpha_i - \mu \alpha_j = 0$ and $\sum_{i=1}^m \lambda_i b_i - \mu b_j < 0$. If μ happens to be zero, we are again violating one of the inequalities in part (1). If $\mu > 0$ then we can assume without loss of generality that $\mu = 1$ and we are violating the inequality $\sum_{i=1}^m \lambda_i b_i \geq b_j$ in part (2).

This proves that the inequalities of types (1) and (2) indeed define the cone $\mathcal{P}(\mathcal{V})$. That the ones with minimal support are enough is easy. Suppose, for example, that we have an inequality of type (1) for which the corresponding nonnegative linear dependence does not have minimal support. Thus, there is another nonnegative dependence $\sum_{i=1}^m \mu_i \alpha_i = 0$ with strictly smaller support. Let $t > 0$ be the greatest number such that the vector $\lambda_i - t\mu_i \geq 0$ for every $i \in \{1, \dots, m\}$. Then, $\sum_{i=1}^m (\lambda_i - t\mu_i) \alpha_i = 0$ is another nonnegative dependence with support strictly contained in that of the original one, and the inequalities for the dependences μ and $\alpha - t\mu$ imply the original inequality for λ . The proof for the inequalities of the second type is similar. \square

Remark 3.25. We need to allow for b_i 's that produce polytopes with fewer than m facets (implying that some of the inequalities in (2) become redundant) in order for $\mathcal{P}(\mathcal{V})$ to be a closed cone. For example, the b_i 's defined by $b_i = \langle \alpha_i, p \rangle$ for a fixed $p \in \mathbb{R}^n$ produce, as mentioned above, the 0-dimensional polytope $\{p\}$. The vectors b of this form are a linear subspace of dimension n of $\mathbb{R}^{\mathcal{V}}$, namely the image of the linear map $\mathbb{R}^n \rightarrow \mathbb{R}^{\mathcal{V}}$ whose matrix has $\alpha_1, \dots, \alpha_m$ as columns, and this linear subspace is the lineality space of $\mathcal{P}(\mathcal{V})$.

If $b, b' \in \mathcal{P}(\mathcal{V})$ then the polytope produced by $b + b'$ is the Minkowski sum of the polytopes of b and of b' . In particular, if one of them, say b is in the lineality space, adding b to a b' corresponds to translating the polytope. That is, modding out the lineality space of $\mathcal{P}(\mathcal{V})$ considered as a cone in $\mathbb{R}^{\mathcal{V}}$ corresponds to modding out translations from $\mathcal{P}(\mathcal{V})$ considered as a space of polytopes.

Example 3.26. Let \mathcal{V} be the set of rays in the Hirzebruch fan, that is, \mathcal{V} consists of

$$\alpha_1 = (0, -1), \alpha_2 = (0, 1), \alpha_3 = (-1, 0), \alpha_4 = (1, k),$$

where k is a positive integer. See Figure 2. We are looking at the polytopes of the form

$$-x_2 \leq b_1, \quad x_2 \leq b_2, \quad -x_1 \leq b_3, \quad x_1 + kx_2 \leq b_4,$$

for different choices of (b_1, b_2, b_3, b_4) .

According to Theorem 3.24 the inequalities that define the cone $\mathcal{P}(\mathcal{V})$ come from the following linear equations among the α_i :

$$\alpha_1 + \alpha_2 = 0, \quad k\alpha_1 + \alpha_3 + \alpha_4 = 0, \quad \alpha_3 + \alpha_4 = k\alpha_2,$$

producing according to Proposition 3.11 the inequalities

$$b_1 + b_2 \geq 0, \quad kb_1 + b_3 + b_4 \geq 0, \quad b_3 + b_4 \geq kb_2.$$

The first inequality is easy to understand. Since we have $-b_1 \leq x_2 \leq b_2$ among the inequalities defining our polytopes, we need $-b_1 \leq b_2$ for the polytope to be non-empty. The second inequality is redundant since it is derived adding to the third k times the first. Hence we have

$$\mathcal{P}(\mathcal{V}) = \left\{ (b_1, b_2, b_3, b_4) \in \mathbb{R}^4 \mid b_1 + b_2 \geq 0, b_3 + b_4 \geq kb_2 \right\}.$$

Let us relate this with the description of Hirzebruch trapezoids in Corollary 2.23. Observe that, as said above, $\mathcal{P}(\mathcal{V})$ has a *lineality space* of dimension two, namely

$$\left\{ (b_1, b_2, b_3, b_4) \in \mathbb{R}^4 \mid b_1 + b_2 = 0, kb_2 + b_3 + b_4 = 0 \right\} = \langle (0, 0, 1, -1), (1, -1, 0, k) \rangle.$$

Modding out this subspace is equivalent to modding out translations in our space of polytopes. There are different ways of doing it, one of which is taking $b_1 = b_2$ and $b_3 = 0$. This corresponds to having the origin as the mid-point of the left edge of the trapezoid, as we did in Corollary 2.23. Once this is done, the parameters a and b from that description correspond to our b_4 and $2b_1 = 2b_2$. Hence, the inequality $b_3 + b_4 \geq kb_2$ that we have now corresponds exactly to the inequality $2a < bk$ that we had there, except for the fact that we are now including the equality case in the cone. This equality case corresponds to the trapezoid degenerating to a triangle. There is another degeneration associated to the inequality $b_1 + b_2 \geq 0$: when it is met with equality the two parallel edges of the trapezoid coincide so the trapezoid becomes a segment. When both inequalities become equalities we are in the lineality space, corresponding to the trapezoid degenerating to a point.

The cone $\mathcal{P}(\mathcal{V})$ gets stratified into a finite number of subcones corresponding to the possible complete polytopal fans that we can obtain using the rays \mathcal{V} .

Definition 3.27. Let \mathcal{N} be a complete fan with set of ray generators contained in $\mathcal{V} = \{\alpha_1, \dots, \alpha_m\} \subset (\mathbb{R}^n)^*$. We call *secondary cone of \mathcal{N}* the closure of

$$\left\{ b \in \mathbb{R}^{\mathcal{V}} \mid \text{the polytope } \bigcap_{i=1}^m \{x \in \mathbb{R}^n \mid \langle \alpha_i, x \rangle \leq b_i\} \text{ has normal fan } \mathcal{N} \right\} \subset \mathcal{P}(\mathcal{V}) \subset \mathbb{R}^{\mathcal{V}}.$$

We show below that the stratification of $\mathcal{P}(\mathcal{V})$ into the secondary cones of the finitely many possible fans with rays in \mathcal{V} is itself a fan called the *secondary fan* of \mathcal{V} . This stratification is quite well understood in geometric combinatorics and is one of the main objects in the book [12] (see, e.g., Theorem 9.5.6 in [12]). Observe that in the language of [12] the complete fans with rays contained in \mathcal{V} are called *polyhedral subdivisions* of \mathcal{V} , with polytopal ones being called *regular subdivisions* and simplicial ones called *triangulations*.

This stratification also plays a role in toric geometry. If \mathcal{N} is a fan with set of rays \mathcal{V} then $\mathbb{R}^{\mathcal{V}}$ can naturally be identified with the space of torus-invariant Weyl divisors, tensored with \mathbb{R} , on the toric variety $X_{\mathcal{N}}$. Under this identification, the (open) secondary cone of \mathcal{N} is the cone spanned by *ample Cartier* divisors in $X_{\mathcal{N}}$ and its closure is spanned by *nef* divisors. Because of this the secondary cone of \mathcal{N} is called the *nef cone* of $X_{\mathcal{N}}$ in [8, Chapter 6].

In particular, the following statement is a rewriting of [12, Corollary 5.2.7]. Its counterpart in toric geometry is [8, Theorem 4.2.8 and Lemma 6.1.13]:

Theorem 3.28. *Let \mathcal{N} be a full-dimensional simplicial fan in $(\mathbb{R}^n)^*$ with set of ray generators $\mathcal{V} = \{\alpha_1, \dots, \alpha_m\}$. Then the secondary cone of \mathcal{N} is a relatively open polyhedral cone in \mathbb{R}^m defined by the following inequalities:*

- (1) *For each linear dependence $\sum_{i=1}^m \lambda_i \alpha_i = 0$ with support contained in a cone of \mathcal{N} , the equality $\sum_{i=1}^m \lambda_i b_i = 0$.*
- (2) *For each linear combination $\sum_{i=1}^m \lambda_i \alpha_i$ with support contained in a cone of \mathcal{N} and producing as a result a certain α_j , the inequality $\sum_{i=1}^m \lambda_i b_i < b_j$.*

Moreover, only the linear dependences and combinations with minimal support (which are finitely many) are needed. The secondary cone is obtained by turning the strict inequalities in part (2) into equalities.

Proof. This is just a rewriting of [12, Corollary 5.2.7]. □

In Example 3.26 the secondary cone of the Hirzebruch trapezoid is the whole cone $\mathcal{P}(\mathcal{V})$, since the Hirzebruch fan is the only simplicial fan using all the rays of \mathcal{V} . We leave it to the reader to check that the systems of inequalities of Theorems 3.24 and Theorems 3.28 are in this case equivalent (when in the second one the inequalities are taken nonstrict). In fact, the same is true for every two dimensional example, since a two-dimensional complete fan is fully characterized by the rays it uses.

Remark 3.29. It must be noted that our treatment of redundant inequalities is different from the one in [12]. Here we insist on those values to equal the support function of the polytope $P(\mathcal{V}, b)$, so in particular the b_i of redundant inequalities are determined by the rest. There, the values b_i corresponding to redundant inequalities are not determined (as long as the corresponding inequalities stay redundant). This means that what we call secondary fan is actually a subfan of the one in [12]. For us only the (polytopal) simplicial fans using *all* the vectors from \mathcal{V} as rays give full-dimensional secondary cones, while in [12] all polytopal simplicial fans with rays contained in \mathcal{V} give full-dimensional cones.

This can be seen in that the analogue in [12] of our Theorem 3.24 is [12, Theorem 4.1.39], which has only the inequalities of type (1) and hence defines a larger cone.

The following summarizes the main results about secondary fans and secondary cones. We refer to [12] for the proofs.

Theorem 3.30 ([12, Proposition 5.2.9 and Theorem 5.2.11]). *Let \mathcal{V} be as above:*

- (1) *The secondary cones $\mathcal{P}(\mathcal{N})$ for the different complete fans \mathcal{N} with rays contained in \mathcal{V} form a fan covering $\mathcal{P}(\mathcal{V})$.*

- (2) The face poset of this fan is isomorphic to the refinement poset of fans. That is, \mathcal{N} refines \mathcal{N}' if and only if the secondary cone of \mathcal{N}' is a face of that of \mathcal{N} .
- (3) The secondary cones of non-simplicial or non-polytopal fans are never full-dimensional.
- (4) The secondary cone of a polytopal simplicial fan with m' rays has dimension m' . In particular, full-dimensional secondary cones correspond bijectively to simplicial polytopal fans with set of rays exactly \mathcal{V} .

Theorem 3.12 implies that the $\|\cdot\|_\infty$ metric that $\mathcal{P}(\mathcal{V})$ inherits from $\mathbb{R}^{\mathcal{V}}$ is equivalent to the Hausdorff metric. More precisely, they are the same metric, modulo rescaling each coordinate i of $\mathbb{R}^{\mathcal{V}}$ by the norm of the corresponding α_i . Hence:

Theorem 3.31. *Let $\mathcal{V} \subset (\mathbb{R}^n)^*$ be a configuration of m non-zero vectors positively spanning $(\mathbb{R}^n)^*$. Then $\mathcal{P}(\mathcal{V})$, considered as a closed polyhedral cone in \mathbb{R}^m with the description of Theorem 3.24, is isometric (modulo rescaling of coordinates) to $\mathcal{P}(\mathcal{V})$ considered as a subspace of $(\mathcal{P}(n), \delta^H)$.*

Corollary 3.32. *Let \mathcal{N} be a complete simplicial fan in $(\mathbb{R}^n)^*$, with m generators. Then, the subset of $\mathcal{P}(n)$ consisting of polytopes whose normal fan equals \mathcal{N} is homeomorphic to \mathbb{R}^m . More precisely, it is isometric to the interior of a full-dimensional polyhedral cone in \mathbb{R}^m with a lineality space of dimension n , where \mathbb{R}^m is considered with the $\|\cdot\|_\infty$ metric.*

3.3.2. *Hausdorff dimension of the space of Delzant polytopes.* We can now answer Problems 5.3 and 5.5 from Fujita-Ohashi [22]. For $m, n \in \mathbb{N}$ let $\mathcal{D}_f^m(n)$ and $\widetilde{\mathcal{D}}_f^m(n)$ denote the spaces of Delzant n -polytopes with m facets, with and without $\text{AGL}(n, \mathbb{Z})$ equivalence.

Theorem 3.33.

- (1) The Hausdorff dimension of both $\mathcal{D}_f^m(n)$ and $\widetilde{\mathcal{D}}_f^m(n)$ is exactly m .
- (2) The Hausdorff dimension of $\mathcal{D}(n)$ and $\widetilde{\mathcal{D}}(n)$ is ∞ .

Proof. Both statements follow easily from the fact that the Hausdorff dimension of a countable union of sets equals the supremum of the Hausdorff dimensions of the individual sets.

For part (1), $\mathcal{D}^m(n)$ is the union of the sets from Corollary 3.32 for the different unimodular normal fans with m rays. There is a countable amount of such fans, and the stratum for each fan has dimension m . In the case of $\widetilde{\mathcal{D}}_f^m(n)$ the stratum of a certain fan \mathcal{N} is the quotient of its stratum in $\mathcal{D}_f^m(n)$ by the finite group of $\text{GL}(n, \mathbb{Z})$ -automorphisms of \mathcal{N} , so the quotient does not change dimension.

For part (2) simply observe that $\mathcal{D}(n)$ and $\widetilde{\mathcal{D}}(n)$ are the countable unions of $\mathcal{D}_f^m(n)$ and $\widetilde{\mathcal{D}}_f^m(n)$ with $m \in N$. \square

Fujita-Ohashi [22, Problems 5.3, 5.5] ask about the dimension of the spaces $\mathcal{D}_v^\ell(n)$ and $\widetilde{\mathcal{D}}_v^\ell(n)$ of Delzant n polytopes with ℓ vertices, rather than facets. In order to translate Theorem 3.33 to this setting we only need to know what the maximum number of facets of a Delzant n -polytope with ℓ vertices is.

Lemma 3.34 (Lower bound theorem for simple polytopes [6] (see also, e.g., [60, Corollary 8.38])). *Let P be a simple n -polytope with m facets. Then, P has at least*

$$(n+1) + (m-n-1)(n-1) = (m-n)(n-1) + 2$$

vertices, with equality if, and only if, the normal fan of P is obtained from that of an n -simplex by $m-n-1$ stellar subdivisions at full-dimensional cones.

Corollary 3.35. *The Hausdorff dimension of both $\mathcal{D}_v^\ell(n)$ and $\widetilde{\mathcal{D}_v^\ell(n)}$ is at most*

$$n + \frac{\ell - 2}{n - 1},$$

with equality if and only if $\ell = 2 \pmod{n - 1}$.

Proof. By the argument in the proof of Theorem 3.33 the Hausdorff dimension equals the maximum number of facets in a Delzant polytope with m vertices. By Lemma 3.34 this is at most $n + \frac{\ell - 2}{n - 1}$ with equality only for iterated stellar subdivisions of a simplex, which necessarily have $\ell = 2 \pmod{n - 1}$ and include iterated corner choppings of Delzant simplices. \square

Remark 3.36. In dimension two $\mathcal{D}_f^m(2) = \mathcal{D}_v^\ell(2)$ is connected, except for the Hirzebruch case $m = 4$ and the case $m = 3$ (the latter is a point in $\widetilde{\mathcal{D}(2)}$ and a discrete space in $\mathcal{D}(2)$). In dimension three or higher we expect every (or infinitely many) of the $\mathcal{D}^m(n)$ to be not connected.

3.4. The fundamental group of $\mathcal{D}(n)$. As mentioned above, the space $\mathcal{P}(n)$ of all polytopes is contractible, since we can homotope the identity map to the constant map $\mathcal{P}(n) \rightarrow P$ by Minkowski segments, where P is any n -polytope. The same is true for $\mathcal{P}_\mathbb{Q}(n)$ but not, at least not obviously, for Delzant polytopes. In what follows we look at the fundamental group of $\mathcal{D}(n)$:

Suppose that a path $\gamma : [0, 1] \rightarrow \mathcal{D}(n)$ has *finite complexity*; that is, the number of rays (equivalently, the number of normal fans) that arise along the path is finite. Then, by Corollary 2.16, there is a polytope Q whose normal fan refines that of $\gamma(t)$ for every t and by Lemma 3.17 we have the following well-defined homotopy:

$$\begin{aligned} \Gamma : [0, 1] \times [0, 1] &\longrightarrow \mathcal{D}(n) \\ (t, s) &\longmapsto (1 - s)\gamma(t) + sQ. \end{aligned}$$

Corollary 3.37. *Loops of finite complexity in $\mathcal{D}(n)$ are null-homotopic.*

Proof. If γ is a loop of finite complexity with $\gamma(0) = \gamma(1) = P$, the homotopy Γ above shows that γ is homotopic to the Minkowski path $PQQP$, which is obviously null-homotopic. \square

A similar argument can be applied to the following class of paths.

Definition 3.38. We say that a path $\gamma : [0, 1] \rightarrow \mathcal{D}(n)$ is *locally refining* if for every $t_0 \in [0, 1]$ there is an ϵ such that the fan of every polytope $\gamma(t)$ with $|t - t_0| \leq \epsilon$ refines that of $\gamma(t_0)$.

Theorem 3.39. *Locally refining loops in $\mathcal{D}(n)$ are null-homotopic.*

Proof. By compactness, there are $0 = t_0 < t_1 < \dots < t_k = 1$ such that each interval $[t_{i-1}, t_i]$, $i \in \{1, \dots, k\}$, contains a point r_i such that the normal fan of $\gamma(t)$ refines that of $\gamma(r_i)$, for every $t \in [t_{i-1}, t_i]$. This makes the following homotopy to be well-defined in $\mathcal{D}(n)$:

$$\begin{aligned} \Gamma : [t_{i-1}, t_i] \times [0, 1] &\longrightarrow \mathcal{D}(n) \\ (t, s) &\longmapsto (1 - s)\gamma(t) + s\gamma(r_i). \end{aligned}$$

Hence, each subpath $\gamma|_{[t_{i-1}, t_i]}$ is homotopic to the Minkowski path $\gamma(t_{i-1})\gamma(r_i)\gamma(t_i)$ and the whole loop γ is homotopic to a loop of finite complexity, which is null-homotopic by Corollary 3.37. \square

Locally refining loops are a more general class than loops of bounded complexity, but they do not include all loops:

Lemma 3.40. *Every loop of finite complexity is locally refining.*

Proof. Suppose that a loop γ has finite complexity, so that along the loop only a finite number of normal fans, say $\mathcal{N}_1, \dots, \mathcal{N}_k$ arise. Observe that this does not imply that the fan changes only at a finite number of values of t , since we can alternate infinitely many times between two or more fans.

Suppose, to seek a contradiction, that γ is not locally refining. That is, there is a $t_0 \in [0, 1]$ such that in every interval $(t_0 - \varepsilon, t_0 + \varepsilon)$ a fan that does not refine that of $\gamma(t_0)$ arises. Call the latter \mathcal{N}_0 . By finiteness of the number of fans, there is a particular fan \mathcal{N} not refining \mathcal{N}_0 and a sequence $(t_i)_{i \in \mathbb{N}}$ converging to t_0 such that the normal fan of $\gamma(t_i)$ equals \mathcal{N} for every t_i . Let \mathcal{V} be the set of primitive generators of rays in \mathcal{N} . Then, the secondary cone of \mathcal{N} in the cone $\mathcal{P}(\mathcal{V})$ of Section 3.3, which is closed in $\mathcal{P}(n)$ since it is complete, contains $\gamma(t_0)$ in its boundary. This contradicts the fact that the boundary of the secondary cone of \mathcal{N} consists of the secondary cones of fans refined by \mathcal{N} , while \mathcal{N}_0 is not refined by \mathcal{N} . \square

Example 3.41 (A loop that is not locally refining). Let

$$P = \text{Conv} \{(0, 0, 0), (1, 0, 0), (0, 1, 0), (0, 0, 1)\}$$

be the standard tetrahedron in \mathbb{R}^3 . Consider the following perturbations of it, with $k \in \mathbb{N}$:

$$P_k = \text{Conv} \{(0, 0, 0), (1, 1/(k+1), -1/k), (0, 1, 0), (0, 0, 1)\}$$

Since the Hausdorff distance between two polytopes is attained at a vertex of one of them, and P and P_k have all but one vertices in common, we have that

$$\delta^H(P, P_k) \leq \|(1, 0, 0) - (1, 1/(k+1), -1/k)\| < 2/k.$$

Of course, P_k is not a Delzant polytope. The crucial property that we need is that its normal fan contains the 2-dimensional cone C_k generated by the vectors

$$(-1, 0, -k) \text{ and } (1, -k-1, 0).$$

This cone has multiplicity one for every k since it is a face of a unimodular cone:

$$\begin{vmatrix} -1 & 0 & -k \\ 1 & -k-1 & 0 \\ 1 & 1 & k+1 \end{vmatrix} = 1$$

The relative interior of C_k intersects the relative interior of $\text{Cone}\{(0, -1, 0), (0, 0, -1)\}$ at the ray generated by $(0, -k-1, -k)$. In particular, no fan containing C_k as a cone refines the normal fan of the standard tetrahedron P .

Now, let Q_k be a Delzant polytope whose normal fan refines that of P_k , obtained by the method in Theorem 2.15; that is, by a sequence of stellar subdivisions at rays lying in the relative interior of cones of multiplicity greater than one. Then, the polytope Q_k still contains the cone C_k in its normal fan, so the normal fan of Q_k does not refine that of the standard simplex P . We also assume that

$$\delta^H(P, Q_k) < 2/k,$$

which is no loss of generality since $\delta^H(P, P_k) < 2/k$ and the Minkowski segment between any Q_k and P_k gives us polytopes with the same normal fan as Q_k and arbitrarily close to P_k .

Finally, let R_k be an arbitrary Delzant polytope whose normal fan refines that of both Q_k and Q_{k-1} , and assumed to still be at distance at most $2/k$ from both P .

We are now ready to define the path γ . Restricted to each interval $[1/k, 1/(k-1)]$ we define γ as the Minkowski path going from Q_k to R_k and then from R_k to Q_{k-1} . The concatenation of this infinitely many paths defines a path $\gamma : (0, 1] \rightarrow \mathcal{D}(3)$, and the fact that all the polytopes in the k -th interval are at distance $\leq 2/k$ from P implies that we can extend this path to $[0, 1]$ by simply defining $\gamma(0) = P$. The path so constructed is not locally refining since in every interval $[0, \varepsilon)$ infinitely many of the normal fans of Q_k appear, and none of them refines that of $P = \gamma(0)$. Concatenating γ with any Minkowski path from $\gamma(1)$ to P we get a loop that is not locally refining.

Remark 3.42. The path γ constructed in Example 3.41 shows that Theorem 3.39 is not enough to conclude the simple connectedness of $\mathcal{D}(n)$ but the path itself is null-homotopic by the following argument. Let R'_k be a fan refining that of P and of R_k , and still assumed to be at distance strictly less than $2/k$ of P . We can homotope γ to the path γ' that uses R'_k instead of R_k in each interval $[1/k, 1/(k-1)]$, and then homotope γ' to a path γ'' ; this substitutes the polytope Q_k that we now have between R'_k and R'_{k+1} by a polytope Q'_k whose fan refines that of R'_k and R'_{k+1} , hence that of P . This produces a path homotopic to γ but in which all normal fans refine that of P .

However, for $n = 2$ we have the following result, stronger than [21, Theorem 5.1.5]:

Theorem 3.43. *If a sequence $(P_k)_{k \in \mathbb{N}}$ of Delzant polygons converges to a Delzant polygon P then there is a k_0 such that the normal fan of P_k refines that of P for every $k \geq k_0$.*

Proof. Since $\mathcal{N}(P)$ has finitely many rays it is enough to concentrate on a particular normal ray α_0 of P and show that there is a k_0 such that every $\mathcal{N}(P_k)$ with $k \geq k_0$ contains that ray. Without loss of generality, let $\alpha_0 = (1, 0)$.

By [21, Corollary 5.1.4] there must be a ray α_k in each $\mathcal{N}(P_k)$ such that the sequence $(\alpha_k)_k = (a_k, b_k)_k$ converges to α_0 . (Here convergence is meant as rays, not as generators; that is, their slopes b_k/a_k converge to zero). This is because P_k needs to have an edge, or a path of them, converging to the edge of P having $\alpha_0 = (1, 0)$ as normal vector.

This α_k may have positive or negative slope, but we can take a subsequence where the slopes have the same sign, and assume it positive. That is, the ray has $a_k > 0$, $b_k \geq 0$ and b_k/a_k converging to zero. Moreover, there is no loss of generality in assuming that α_k is the last ray in $\mathcal{N}(P_k)$, in clockwise order, with these properties, since any ray between α_k and the ray $(1, 0)$ would have a smaller slope and would still give a sequence with slopes converging to zero.

Then, the ray of $\mathcal{N}(P_k)$ after α_k in clockwise direction is of the form $\beta_k = (-c_k, -d_k)$ with $d_k > 0$. For the corresponding cone to be unimodular we need $a_k d_k - c_k b_k = 1$, which implies that c_k is also positive and that c_k/d_k converges to infinity too. That is, the rays β_k lie in the negative quadrant and have slopes converging to zero.

Now, let $\beta = (e, -1)$ be the next ray after $(1, 0)$ in the normal fan of P , in clockwise order (e can be positive or negative). The contradiction is that the polygons P_k need to also have rays converging to β , but this is impossible by the presence of the cone $\text{Cone}\{\alpha_k, \beta_k\}$ in $\mathcal{N}(P_k)$. \square

Corollary 3.44. *All paths in $\mathcal{D}(2)$ are locally refining.*

Proof. By Theorem 3.43, each Delzant polygon P has a neighborhood in $\mathcal{D}(2)$ such that the normal fan of every polygon in the neighborhood refines that of P . \square

Corollary 3.45. *$\mathcal{D}(2)$ is simply connected.*

Proof. Every loop in $\mathcal{D}(2)$ is locally refining by the Corollary 3.44, and locally refining loops are null-homotopic by Theorem 3.39. \square

Question 3.46. *Is the fundamental group of $\mathcal{D}(n)$ trivial?*

3.5. The CW topology on $\mathcal{D}(n)$. Since $\mathcal{P}(n)$ is the union of the subsets $\mathcal{P}(\mathcal{N})$ for the different complete fans \mathcal{N} and since each $\mathcal{P}(\mathcal{N})$ is isometric to a finite dimensional polyhedral cone, the stratification by normal equivalence is a cell decomposition of $\mathcal{P}(n)$ into the uncountably many strata $\mathcal{P}(\mathcal{N})$ for the possible fans. The closure of each cell is isometric to a closed polyhedral cone in some \mathbb{R}^m with its face structure being the poset structure of the finitely many fans refined by \mathcal{N} . In order to make these closures compact, we only need to take the quotient of $\mathcal{P}(n)$ by translations and positive scaling. Indeed, the quotient by translations removes the common lineality space of all closed cells, turning them into pointed cones, and the quotient by positive scaling makes these pointed cones compact; indeed, homeomorphic to polytopes, and preserving their face structure. That is:

Proposition 3.47. *The quotient of $\mathcal{P}(n)$ by translations and positive scaling has a natural structure of regular CW complex in which closed cells of dimension m correspond to secondary cones of dimension $m + n + 1$ and the face poset structure of the CW complex is that of refinement of normal fans.*

In this structure $\mathcal{P}_{\mathbb{Q}}(n)$ is a subcomplex with countably many cells, and $\mathcal{D}(n)$ is a union of open cells, but not a subcomplex.

We call the topology induced on any of $\mathcal{P}(n)$, $\mathcal{P}_{\mathbb{Q}}(n)$ and $\mathcal{D}(n)$ spaces by the CW structure their CW *topology*. Observe that a priori this is defined on a quotient space, but it can be lifted to the original space via the following considerations:

- One canonical way to choose a representative from each class in the quotient is to translate and dilate every polytope other than a single point (a 0-polytope) so that its circumball is the unit ball. The *circumball* or smallest enclosing ball exists and is unique for every bounded subset of \mathbb{R}^n [24, p. 49]. This shows that the quotient space is also a subspace.
- Let $\widetilde{\mathcal{P}}(n)$ denote the quotient of $\mathcal{P}(n) \setminus \{0\text{-polytopes}\}$ by translations and dilations. The following map is a bijection:

$$\begin{aligned} \mathcal{P}(n) \setminus \{0\text{-polytopes}\} &\longrightarrow \widetilde{\mathcal{P}}(n) \times \mathbb{R}^n \times (0, \infty) \\ P &\longmapsto (\tilde{P}, c(P), r(P)) \end{aligned}$$

where \tilde{P} , $c(P)$ and $r(P)$ denote, respectively, the class of P , the center, and the radius of the circumball of P .

Via this bijection, we can lift the CW topology of $\widetilde{\mathcal{P}}(n)$ to $\mathcal{P}(n)$, and to its subspaces $\mathcal{P}_{\mathbb{Q}}(n)$ and $\mathcal{D}(n)$. Since $\mathbb{R}^n \times (0, \infty)$ is contractible, the homotopy type of $\mathcal{D}(n)$ and its quotient is the same.

Regarded from this perspective, the fans constructed in Theorem 2.27 correspond to open cells that are in $\mathcal{D}(n)$ but with empty boundary in $\mathcal{D}(n)$. Also, this provides another way of expressing why connectivity by blow-ups/blow-downs as we have for $n = 2$ is nicer than the one we have for $n \geq 3$: for $n = 2$ there are paths that move only through cells of consecutive dimensions.

However, it must be noted that the CW topology obtained in this way is not the metric topology of $\mathcal{P}(n)$ that we are considering in the paper, coming from δ^H or δ^V . In fact, we have the following:

Lemma 3.48. *The CW topology on $\mathcal{P}(n)$ (or on $\mathcal{P}_{\mathbb{Q}}(n)$, or on $\mathcal{D}(n)$) is not metrizable. Hence, it does not coincide with the topology coming from $\delta^V \circ \delta^H$.*

Proof. A metrizable CW complex must be locally finite, i.e., every point must lie only in finitely many closed cells. In our case every polytope P lies in the infinitely many cells corresponding to fans that refine $\mathcal{N}(P)$. \square

Proposition 3.49. *The CW topology on $\mathcal{P}(n)$ (or on $\mathcal{P}_{\mathbb{Q}}(n)$, or on $\mathcal{D}(n)$) is strictly finer than the metric topology coming from $\delta^V \circ \delta^H$.*

Proof. In the CW complex a set is, by definition, open (respectively, closed) if and only if its intersection with each closed cell is open (respectively, closed). Since the closed cells of the CW complex are closed also in the metric topology, the CW topology is finer. It is strictly finer since, for example, $\mathcal{P}_{\mathbb{Q}}(n)$ is closed in $\mathcal{P}(n)$ with respect to the CW topology (as is any union of closed cells in a CW complex) while it is dense in the metric topology (cf. Theorem 3.18). \square

However, for every finite, or even locally finite, subcomplex the CW topology and the metric topology in $\mathcal{P}(n)$ coincide.

As another proof that the topologies are different, Example 3.41 shows a path γ in $\mathcal{D}(3)$ with respect to the metric topology in which every neighborhood of $P_0 = \gamma(0)$ contains polytopes whose fan does not refine that of γ_0 . In the CW topology, the union of cells corresponding to polytopes that refine $\mathcal{N}(P)$ are, by definition, a neighborhood of P . More strongly, in the CW topology all paths have finite complexity:

Lemma 3.50. *Let $n \in \mathbb{N}$ and let X be a compact topological space. A map $f : X \rightarrow \mathcal{D}(n)$ is continuous with respect to the CW topology in $\mathcal{D}(n)$ if and only if:*

- (1) *It is continuous with respect to the metric topology, and*
- (2) *$f(X)$ intersects only finitely many cells.*

Proof. Necessity of (1) follows from Proposition 3.49 and necessity of (2) from the fact that $f(X)$ is compact and every compact subspace in a CW complex is contained in a finite subcomplex [28, Proposition A.1].

In order to prove sufficiency, suppose that f intersects only k cells, corresponding to the unimodular fans $\Sigma_1, \dots, \Sigma_k$. Let Σ be a common unimodular refinement of them and let \mathcal{V} the finite set of rays used in Σ . Then f can be considered as a map $X \rightarrow f(X) \subset \mathcal{P}(\mathcal{V})$ and in $\mathcal{P}(\mathcal{V})$ the CW topology and the metric topology coincide. \square

Recall that a space with all homotopy groups trivial is called *weakly contractible*.

Corollary 3.51. *$\mathcal{D}(n)$ with its CW topology is weakly contractible.*

Proof. Let $f : \mathbb{S}^k \rightarrow \mathcal{D}(n)$ be a continuous map from the k -sphere \mathbb{S} to $\mathcal{D}(n)$. We need to show that it is homotopic to a constant map. By [28, p. 346] to show that a homotopy group is trivial we can forget basepoints. By Lemma 3.50 there is a Delzant polytope Q whose normal fan refines $f(x)$ for every $x \in \mathbb{S}^k$. We consider the following homotopy, similar to the one in Corollary 3.37 and Theorem 3.39:

$$\begin{aligned} \Gamma : \mathbb{S}^k \times [0, 1] &\longrightarrow \mathcal{D}(n) \\ (x, s) &\longmapsto (1 - s)f(x) + sQ. \end{aligned}$$

This map is continuous with respect to the metric topology and its image intersects only the cells which $f(\mathbb{S}^k)$ intersects plus the cell corresponding to Q . \square

Remark 3.52. A *weak homotopy equivalence* between two topological spaces is a continuous map inducing isomorphisms in all homotopy groups. Whitehead’s Theorem [28, Theorem 4.5] says that a weak homotopy equivalence between CW complexes is in fact a homotopy equivalence, so weakly contractible CW complexes are contractible. We cannot apply this directly to $\mathcal{D}(n)$ since it is a union of open cells, but not a subcomplex, of $\mathcal{P}(n)$.

4. TRANSLATION TO SYMPLECTIC TORIC MANIFOLDS

In this section we translate the main result from Sections 2 and 3 into the language of symplectic toric manifolds.

4.1. Symplectic toric manifolds: basic theory. For an introduction to symplectic geometry and Hamiltonian torus actions we refer to Cannas da Silva [10], Guillemin-Sjamaar [26] and Pelayo [49]. The origins of symplectic geometry can be traced back to classical mechanics, see Abraham-Marsden [1], De León-Rodríguez [11] and Marsden-Ratiu [41] for texts with an emphasis on this point of view.

4.1.1. Definition of symplectic toric manifold. We start reviewing the concept of symplectic manifolds, which is the mathematical model of the phase space of physics.

Definition 4.1. A *symplectic manifold* (M, ω) is a pair consisting of a smooth manifold M and a *symplectic form* ω , i.e., a non-degenerate closed 2-form on M . We denote by $\text{Symp}(M)$ the group of symplectomorphisms (i.e. symplectic diffeomorphisms) of M .

Throughout this section we write

$$\mathbb{T}^n := (\mathbb{T}^1)^n = (\mathbb{R}/\mathbb{Z})^n$$

for the n -dimensional standard torus. Denote by \mathfrak{t} the Lie algebra $\text{Lie}(\mathbb{T}^n)$ of \mathbb{T}^n and by \mathfrak{t}^* the dual of \mathfrak{t} . Let $\exp: \mathfrak{t} \rightarrow \mathbb{T}^n$ be the exponential map of Lie theory.

Let $\mathbb{T}^n \times M \rightarrow M, (t, m) \mapsto t \cdot m$ be a symplectic action (i.e. an action by symplectomorphisms) of the n -dimensional torus \mathbb{T}^n on (M, ω) . Let $X \in \mathfrak{t}$. Then X generates a smooth vector field X_M on M , called the *infinitesimal generator of the action*, defined by

$$X_M(m) := \left. \frac{d}{dt} \right|_{t=0} \exp(tX) \cdot m.$$

We write $\iota_{X_M}\omega := \omega(X_M, \cdot) \in \Omega^1(M)$ for the contraction 1-form.

Definition 4.2. Let $\langle \cdot, \cdot \rangle: \mathfrak{t}^* \times \mathfrak{t} \rightarrow \mathbb{R}$ be the duality pairing.⁵ The \mathbb{T}^n -action on (M, ω) is said to be *Hamiltonian* if there exists a smooth \mathbb{T}^n -invariant map $\mu: M \rightarrow \mathfrak{t}^*$, called the *momentum map* or *moment map*, such that for all $X \in \mathfrak{t}$ we have $\iota_{X_M}\omega = d\langle \mu, X \rangle$.

In order to simplify the presentation we are going to assume that the momentum map takes values in \mathbb{R}^n , where n is the dimension of \mathfrak{t}^* . This will allow us to work and reason with Delzant polytopes etc in \mathbb{R}^n , as we did in the previous sections, instead of with the abstract space \mathfrak{t}^* . This is a standard simplification (which on the other hand represents no loss of generality) done via the following non-canonical procedure. Let us choose an epimorphism $E: \mathbb{R} \rightarrow \mathbb{T}^1$, for instance, $x \mapsto e^{2\pi\sqrt{-1}x}$. This Lie group epimorphism has discrete center \mathbb{Z} and the inverse of the corresponding Lie algebra isomorphism is given by

⁵In previous sections we used the same notation for the pairing in general real vector spaces or in \mathbb{R}^n .

$\text{Lie}(\mathbb{T}^1) \ni \frac{\partial}{\partial x} \mapsto \frac{1}{2\pi} \in \mathbb{R}$. Therefore, for $\mathbb{T}^n = (\mathbb{T}^1)^n$, we obtain the following isomorphism of Lie algebras

$$\begin{aligned} \text{Lie}(\mathbb{T}^n) = \mathfrak{t} &\longrightarrow \mathbb{R}^n \\ \frac{\partial}{\partial x_k} &\longmapsto \frac{1}{2\pi} e_k, \end{aligned}$$

where e_k is the k th standard basis vector in \mathbb{R}^n . Choosing an inner product on \mathfrak{t} , we obtain an isomorphism $\mathfrak{t} \rightarrow \mathfrak{t}^*$, and hence taking its inverse and composing it with the isomorphism $\mathfrak{t} \rightarrow \mathbb{R}^n$ described above, we get an isomorphism $\mathcal{I} : \mathfrak{t}^* \rightarrow \mathbb{R}^n$. In this way, we obtain a momentum map

$$\mu = \mu_{\mathcal{I}} : M \rightarrow \mathbb{R}^n.$$

Definition 4.3. A *symplectic toric manifold* is quadruple of the form

$$(M, \omega, \mathbb{T}^n, \mu),$$

where (M, ω) is a compact connected symplectic manifold of dimension $2n$ endowed with an effective Hamiltonian action of an n -dimensional torus \mathbb{T}^n admitting a momentum map $\mu : M \rightarrow \mathfrak{t}^*$. Via the (non-canonical) isomorphism $\mathfrak{t}^* \cong \mathbb{R}^n$ the map μ gives rise to a map $M \rightarrow \mathbb{R}^n$ which, to simplify notation, we also denote by $\mu : M \rightarrow \mathbb{R}^n$.

Often symplectic toric manifolds are referred to as (compact connected) *toric integrable systems*. Essentially a toric integrable system on a $2n$ -dimensional symplectic manifold (M, ω) is a smooth map $\mu : (M, \omega) \rightarrow \mathbb{R}^n$ which is an integrable system in the usual sense (see for instance [53] for the basic definitions and properties) in which all components $\mu_i : M \rightarrow \mathbb{R}$ generate periodic flows of the same period. Usually one does not require the manifold to be compact when speaking of toric integrable systems, unlike in the case of symplectic toric manifolds, but other than that they are the same object.

Symplectic toric manifolds also have the structure of a toric variety; see Delzant [14] for this connection and Duistermaat-Pelayo [16] for explicit coordinate formulas explaining it. For surveys concerning symplectic geometry from the point of view of Hamiltonian torus actions and integrable systems, including symplectic toric manifolds, we refer to [53, 49].

4.1.2. Isomorphisms and weak isomorphisms of symplectic toric manifolds.

Definition 4.4. Let $(M, \omega, \mathbb{T}^n, \mu)$ and $(M', \omega', \mathbb{T}^n, \mu')$ be $2n$ -dimensional symplectic toric manifolds endowed with effective symplectic actions $\rho : \mathbb{T}^n \rightarrow \text{Symp}(M, \omega)$ and $\rho' : \mathbb{T}^n \rightarrow \text{Symp}(M', \omega')$. We say that $(M, \omega, \mathbb{T}^n, \mu)$ and $(M', \omega', \mathbb{T}^n, \mu')$ are:

- *isomorphic* if there exists an equivariant symplectomorphism $\varphi : M \rightarrow M'$ such that $\mu' \circ \varphi = \mu$.
- *weakly isomorphic* if there exists an automorphism $h : \mathbb{T}^n \rightarrow \mathbb{T}^n$ and an h -equivariant symplectomorphism $\varphi : M \rightarrow M'$, i.e., the following diagram commutes:

$$\begin{array}{ccc} \mathbb{T}^n \times M & \xrightarrow{\rho^*} & M \\ (h, \varphi) \downarrow & & \varphi \downarrow \\ \mathbb{T}^n \times M' & \xrightarrow{\rho'^*} & M', \end{array}$$

where $\rho^*(x, m) := \rho(x)(m)$ for every $x \in \mathbb{T}^n$, $m \in M$ and ρ'^* is defined analogously.

Since the automorphism group of \mathbb{T}^n is the general linear group $\text{GL}(n, \mathbb{Z})$ we have that h is represented by a matrix $A \in \text{GL}(n, \mathbb{Z})$.

4.1.3. *Atiyah-Guillemin-Sternberg Convexity Theorem.* Let (M, ω) be a symplectic manifold of dimension $2n$ and let $(\mathbb{T}^1)^k$ be a torus of any dimension $k \leq n$ acting on M Hamiltonianly with momentum map $\mu: M \rightarrow \mathbb{R}^k$. A seminal result proven independently by Atiyah [5] and Guillemin-Sternberg [27] says that the momentum map image

$$\mu(M) \subset \mathbb{R}^k$$

is equal to the polytope $P \subset \mathbb{R}^k$ obtained by taking the convex hull of the images of the fixed point set of the $(\mathbb{T}^1)^k$ -action on M . In our case, since we are interested in symplectic toric manifolds, the relevant case is $k = n$. So if $(M, \omega, \mathbb{T}^n, \mu)$ is a $2n$ -dimensional symplectic toric manifold then $\mu(M)$ is a polytope in \mathbb{R}^n .

This general result by Atiyah and Guillemin-Sternberg has some important precedents in the work of Horn [32], Kostant [37] and Schur [57].

4.1.4. *Definition of the spaces of symplectic toric manifolds.*

Definition 4.5 (Sections 1.2 and 1.3 in [50]). Assuming the previous conventions for \mathbb{T}^n and that the identification $\mathcal{I}: \mathfrak{t}^* \rightarrow \mathbb{R}^n$ is *fixed*, we define:

- $\mathcal{M}(2n) := \mathcal{M}(2n)^{\mathcal{I}}$ as the set of equivalence classes of $2n$ -dimensional symplectic toric manifolds under the equivalence relation given by isomorphisms⁶.
- $\widetilde{\mathcal{M}}(2n)$ as the set of equivalence classes of $2n$ -dimensional symplectic toric manifolds under the equivalence relation given by weak isomorphisms⁷.

To simplify the notations we often write $(M, \omega, \mathbb{T}^n, \mu)$ identifying the representative with its equivalence class $[(M, \omega, \mathbb{T}^n, \mu)]$ in $\mathcal{M}(2n)$.

4.1.5. *Delzant correspondence: from symplectic toric manifolds to Delzant polytopes.* We saw in Section 4.1.3 that if $(M, \omega, \mathbb{T}^n, \mu)$ is a $2n$ -dimensional symplectic toric manifold then $\mu(M)$ is a polytope in \mathbb{R}^n .

A famous theorem by Delzant showed that the polytopes that arise this way are exactly the Delzant polytopes (Definition 2.10) and that the Delzant polytope of a symplectic toric manifold is a complete invariant of its isomorphism type. This was complemented by Karshon-Kessler-Pinsounnault [36, Proposition 2.3 (2)] who gave the analogous result for weak isomorphisms:

Theorem 4.6 (Delzant [14], Karshon-Kessler-Pinsounnault [36]). *Let $\mathcal{D}(n)$ and $\widetilde{\mathcal{D}}(n)$ be the space of Delzant n -polytopes and its quotient by $\text{AGL}(n, \mathbb{Z})$ -equivalence, as defined in Section 3.1. Then*

- *The map*

$$(3) \quad \begin{aligned} \mathcal{M}(2n) &\longrightarrow \mathcal{D}(n). \\ [(M, \omega, \mathbb{T}^n, \mu)] &\longmapsto \mu(M). \end{aligned}$$

is a bijection between $\mathcal{M}(2n)$ and $\mathcal{D}(n)$.

- *This bijection descends to a bijection*

$$(4) \quad \widetilde{\mathcal{M}}(2n) \longrightarrow \widetilde{\mathcal{D}}(n).$$

⁶If we choose a different identification $\mathcal{I}': \mathfrak{t}^* \rightarrow \mathbb{R}^n$, then $\mathcal{M}(2n)^{\mathcal{I}'}$ and $\mathcal{M}_{\mathbb{T}^n}$ are in general different as sets. However they are in bijective correspondence by a bijection which preserves the structures with which we are concerned in the present paper.

⁷Two weakly isomorphic toric manifolds are isomorphic if h is the identity and $\mu' \circ \varphi = \mu$.

Example 4.7 (See Corollary 2.23 and Figure 2).

$$P_\lambda = \left\{ (x_1, x_2) \in \mathbb{R}^2 \mid x_1 \geq 0, x_2 \geq 0, x_1 + x_2 \leq \lambda \right\},$$

is the image under this map of the standard \mathbb{T}^2 action on $\mathbb{C}\mathbb{P}^2$ endowed with the Fubini-Study symplectic form multiplied by λ . The *Hirzebruch trapezoid*,

$$H_{a,b,k} = \left\{ (x_1, x_2) \in \mathbb{R}^2 \mid -\frac{b}{2} \leq x_2 \leq \frac{b}{2}, 0 \leq x_1 \leq a - kx_2 \right\},$$

is the image under this map of the standard toric action on a Hirzebruch surface.

4.2. Geometry and topology of the space of symplectic toric manifolds.

4.2.1. *Symplectic interpretation of Theorems 2.26 and 2.27.* Theorem 2.26 and Theorem 2.27 imply that, unlike in dimension 4, there is no collection of “toric minimal models” for symplectic toric manifolds in dimension 6 or higher. We refer to Delzant [14] and Duistermaat-Pelayo [16] for the relation between a symplectic toric manifold and the underlying toric variety.

Theorem 4.8. *Let (M, ω) be a 6-dimensional symplectic toric manifold. Then there exists another 6-dimensional symplectic toric manifold (M', ω') such that:*

- *there is a toric equivariant epimorphism between the underlying toric varieties $M \twoheadrightarrow M'$,*
- *(M', ω') is not equivariantly symplectomorphic to the toric equivariant blow up of any other 6-dimensional symplectic toric manifold.*

Corollary 4.9. *Given any collection \mathcal{F} of 6-dimensional symplectic toric manifolds there exists another 6-dimensional symplectic toric manifold which cannot be obtained by a finite sequence of toric equivariant symplectic blow ups from an element of \mathcal{F} .*

Theorem 4.10. *Let $n \geq 3$. There is an infinite family $(M_k, \omega_k)_{k \in \mathbb{N}}$ of pairwise non-homeomorphic $2n$ -dimensional symplectic toric manifolds with the property that any toric equivariant epimorphism to another $2n$ -dimensional symplectic toric manifold $M_k \twoheadrightarrow M$ is a homeomorphism.*

4.2.2. *Metric and topology on the moduli spaces $\mathcal{M}(2n)$ and $\widetilde{\mathcal{M}}(2n)$.* The bijection (3) of Theorem 4.6 allows us, following [50] to define a metric on $\mathcal{M}(2n)$ as the pullback of the metric δ^V defined on $\mathcal{D}(n)$ (Section 3.1), thereby getting the metric space $(\mathcal{M}(2n), d)$. This metric induces a topology ν on $\mathcal{M}(2n)$ and it makes the map of Theorem 4.6(1) a homeomorphism.

We also endow $\widetilde{\mathcal{M}}(2n)$ with the topology $\widetilde{\nu}$ induced by (4), that is, with the quotient topology coming from $\mathcal{M}(2n)$, which makes Theorem 4.6(2) a homeomorphism too.

The metric δ^V is very closely related to the famous *Duistermaat-Heckman measure* m_{DH} on \mathbb{R}^n from [15]: if $\mu: (M, \omega) \rightarrow \mathfrak{t}^* \simeq \mathbb{R}^n$ is the momentum map of a symplectic toric manifold then the Duistermaat-Heckman measure on $\mathfrak{t}^* \simeq \mathbb{R}^n$ is given as follows: if $U \subset \mathfrak{t}^* \simeq \mathbb{R}^n$ is a Borel set then $m_{\text{DH}}(U)$ is (up to a constant, depending on conventions) the volume of $\mu^{-1}(U)$ with respect to ω^n , that is, the *symplectic volume* $\int_{\mu^{-1}(U)} \frac{\omega^n}{n!}$. Because the toric case is so especial, in terms of the Euclidean volume (Lebesgue measure on \mathbb{R}^n) we have that $m_{\text{DH}}(U) = \text{Vol}(U \cap \mu(M))$ (in general for Hamiltonian actions of tori of lower dimension $m < n$, the expression of m_{DH} is given by the integral of a function which is a polynomial of degree at most $n - m$ on each component of the regular values of μ , this being a famous result by Duistermaat-Heckman).

4.2.3. *Main Theorem describing the geometry and topology of $\widetilde{\mathcal{M}}(2n)$ and $\mathcal{M}(2n)$.* It follows from the work we have described in this section that Theorems 3.18 and 3.19 are equivalent to the following.

Theorem 4.11. *Let $\mathcal{M}(2n)$ and $\widetilde{\mathcal{M}}(2n)$ be the moduli spaces of toric $2n$ -dimensional manifolds, under isomorphisms and equivariant isomorphisms, respectively. Then:*

- (i) $\mathcal{M}(2n)$ and $\widetilde{\mathcal{M}}(2n)$ are path connected.
- (ii) $\mathcal{M}(2n)$ is not locally compact, but it is dense in $\mathcal{C}_p(n)$ in the sense of identifying $\mathcal{M}(2n)$ with $\mathcal{D}(n)$ via (3). Hence, its completion is the same as that of $\mathcal{C}_p(n)$ (namely $\mathcal{C}_p(n) \cup \{0\}$ for δ^V and $\mathcal{C}(n)$ for δ^H).
- (iii) $(\widetilde{\mathcal{M}}(2n), \widetilde{\delta^V})$ is not locally compact, but it is dense in $\widetilde{\mathcal{C}}_p(n)$ in the sense of identifying $\widetilde{\mathcal{M}}(2n)$ with $\widetilde{\mathcal{D}}(n)$ via (4). Hence, its completion is the same as that of $\widetilde{\mathcal{C}}_p(n)$, namely $\widetilde{\mathcal{C}}_p(n) \cup \{0\}$.

Theorem 4.11 essentially solves a problem proposed by the first author and Vũ Ngọc [54, Problem 2.42] in 2012.

Similarly, Corollary 3.45 implies:

Proposition 4.12. $\mathcal{M}(4)$ is simply connected.

REFERENCES

1. R. Abraham, J.E. Marsden: *Foundation of Mechanics*, Second edition, revised and enlarged. With the assistance of Tudor Ratiu and Richard Cushman. Benjamin/Cummings Publishing Co., Inc., Advanced Book Program, Reading, Mass., 1978.
2. K. Adiprasito, I. Pak: *All triangulations have a common stellar subdivision*, preprint arXiv:2404.05930v2, April 2024, <https://doi.org/10.48550/arXiv.2404.05930>.
3. G. Averkov, J. Hofscheier, B. Nill: Generalized flatness constants, spanning lattice polytopes, and the Gromov width, *Manuscripta Mathematica* volume 170, pages 147–165 (2023).
4. D. Abramovich, K. Karu, K. Matsuki, J. Włodarczyk: Torification and factorization of birational maps *J. Am. Math. Soc.*, 15: 3 (2002), 531–572.
5. M. Atiyah: Convexity and commuting Hamiltonians, *Bull. Lond. Math. Soc.* 14 (1982) 1–15.
6. D.W. Barnette: The minimum number of vertices of a simple polytope, *Israel J. Math.* 10 (1971), 121–125.
7. T. Bogart, C. Haase, M. Hering, B. Lorenz, B. Nill, A. Paffenholz, G. Rote, F. Santos, H. Schenck: Finitely many smooth d -polytopes with N lattice points, *Israel J. Math.* 2015, Volume 207, Issue 1, pp 301–329. DOI:
8. D. Cox, J. Little, H. Schenck: *Toric Varieties*, Graduate Studies in Mathematics, **124**, American Math. Society, 2010.
9. I. Bárány, N. Tokushige: The minimum area of convex lattice n -gons. *Combinatorica*, 24(2):171–185, 2004.
10. A. Cannas da Silva: *Lectures in Symplectic Geometry*, Springer-Verlag, Berlin (2000).
11. M. de León, P. R. Rodrigues: *Methods of Differential Geometry in Analytical Mechanics*, North-Holland Mathematics Studies, 158, North-Holland, Amsterdam, 1989.
12. J. A. De Loera, J. Rambau, F. Santos: *Triangulations: Structures for Algorithms and Applications*, xiii+535 pp. Algorithms and Computation in Mathematics, Vol. 25, Springer-Verlag.
13. C. Delaunay: On Hyperbolicity of Toric Real Threefolds, *IMRN International Mathematics Research Notices* 2005, No. 51.
14. T. Delzant: Hamiltoniens périodiques et images convexes de l’application moment, *Bull. Soc. Math. France* **116**, 315–339 (1988).
15. J.J. Duistermaat, G.J Heckman: On the variation in the cohomology of the symplectic form of the reduced phase space, *Invent. Math.* 69, 259–268 (1982).
16. J.J. Duistermaat, Á. Pelayo: Reduced phase space and toric variety coordinatizations of Delzant spaces, *Math. Proc. Camb. Phil. Soc.* **146**, 695–718 (2009).

17. G. Ewald: *Combinatorial Convexity and Algebraic Geometry*, Graduate Texts in Mathematics 168, Springer-Verlag, 1996.
18. A. Figalli, J. Palmer, Á. Pelayo: Symplectic G-capacities and integrable systems, *Annali della Scuola Normale Superiore di Pisa (5) Vol. XVIII* (2018) 65-103.
19. A. Figalli, Á. Pelayo: On the continuity of ball packing density on moduli spaces of toric manifolds, *Advances in Geometry* 16 (2016) 291-300.
20. O. Fujino and H. Sato: Introduction to the toric Mori theory, *Michigan Math. J.* 52 (2004), 649–665.
21. H. Fujita, Y. Kitabeppu, A. Mitsuishi: Distance functions on convex bodies and symplectic toric manifolds, preprint arXiv:2003.02293, March 2020, <https://doi.org/10.48550/arXiv.2003.02293>.
22. H. Fujita, K. Ohashi: A metric on the moduli space of bodies, preprint arXiv:1804.05161, April 2018, <https://doi.org/10.48550/arXiv.1804.05161>.
23. W. Fulton: *Introduction to Toric Varieties*, Annals of Mathematical Studies, 131. The William H. Roever Lectures in Geometry. Princeton University Press, Princeton, NJ (1993).
24. P. M. Gruber: *Convex and Discrete Geometry*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 336, Springer-Verlag, Berlin, 2007.
25. V. Guillemin: *Moment maps and combinatorial invariants of T^n -spaces*, Birkhäuser, Boston, 1994.
26. V. Guillemin, R. Sjamaar: *Convexity Properties of Hamiltonian Group Actions*. CRM Monograph Series, 26. American Mathematical Society, Providence, RI, 2005. iv+82 pp.
27. V. Guillemin, S. Sternberg: Convexity properties of the moment mapping, *Invent. Math.* 67 (1982) 491-513.
28. A. Hatcher: *Algebraic Topology*, Cambridge University Press, 2002.
29. J.-C. Hausmann, A. Knutson: Polygon spaces and Grassmannian. *L'Enseign. Math.* 43 (1997) 173-198.
30. J.-C. Hausmann, A. Knutson: The cohomology ring of polygon spaces. *Ann. l'Inst. Fourier* 48 (1998) 281-321.
31. H. Hironaka: On the theory of birational blowing-up. *Harvard University Ph. D. Thesis*, 1960.
32. A. Horn: Doubly stochastic matrices and the diagonal of a rotation matrix. *Am. J. Math.* 76 (1954) 620-630.
33. D. Kane, Á. Pelayo, J. Palmer: Classifying toric and semitoric fans by lifting equations from $SL(2, \mathbb{Z})$, *Symmetry, Integrability, and Geometry: Methods and Applications SIGMA* 14 (2018), 016, 43 pages.
34. D. Kane, Á. Pelayo, J. Palmer: Minimal models of compact symplectic semitoric manifolds, *Journal of Geometry and Physics* 125 (2018) 49-74.
35. M. Kapovich, J. Millson: On the moduli space of polygons in the Euclidean plane. *J. Diff. Geom.* 42(1), 133-164 (1995)
36. Y. Karshon, L. Kessler, M. Pinsonnault: A compact symplectic four-manifold admits only finitely many inequivalent toric actions, *J. Symplectic Geom.* 5, 139–166 (2007).
37. B. Kostant: On convexity, the Weyl group and the Iwasawa decomposition, *Ann. Sci. École Norm. Sup. (4)* 6 (1973) 413-455 (1974).
38. M. Krein, D. Milman: On extreme points of regular convex sets, *Studia Math.* 9, 133–138 (1940).
39. A. Lundman: A classification of smooth convex 3-polytopes with at most 16 lattice points. *Journal of Algebraic Combinatorics* 37 (2013), 139–165.
40. A. M. Macbeath: A compactness theorem for affine equivalence-classes of convex regions, *Canad. J. Math.* 3 (1951) 54-61.
41. J.E. Marsden, T.S. Ratiu: *Introduction to Mechanics and Symmetry*, Texts in Applied Mathematics 17, second edition, second printing, Springer Verlag, New York, 2003.
42. D. McDuff and S. Tolman: Polytopes with mass linear functions. I, *International Mathematics Research Notices. IMRN*, 8 (2010), 1506–1574.
43. R. Morelli: The birational geometry of toric varieties. *J. Algebr. Geom.* 5, 751–782 (1996).
44. T. Oda: *Convex Bodies and Algebraic Geometry*, Ergebnisse der Mathematik und ihrer Grenzgebiete, Vol. 15, Springer-Verlag, Berlin, 1988.
45. T. Oda: *Lectures on Torus embeddings and applications*. (Based on joint work with Katsuya Miyake). Bombay: Tata Inst. Fund. Res. 58, Springer, 1978.
46. J. Palmer: Moduli spaces of semitoric systems, *J. Geom. and Physics* 115 (2017) 191-217.
47. Á. Pelayo: Toric symplectic ball packing, *Topology Appl.* 153 (2006) 3633-3644.
48. Á. Pelayo: Topology of spaces of equivariant symplectic embeddings, *Proc. Amer. Math. Soc.* 135 (2007) 277-288.

49. Á. Pelayo: Hamiltonian and symplectic symmetries: An introduction, *Bull. Amer. Math. Soc.* 54 (2017), No. 3, 383-436.
50. Á. Pelayo, A. R. Pires, T. S. Ratiu, S. Sabatini Moduli spaces of toric manifolds, *Geometriae Dedicata* volume 169, pages 323–341 (2014).
51. Á. Pelayo, S. Vũ Ngọc: Semitoric integrable systems on symplectic 4-manifolds, *Invent. Math.* 177 (2009) 571-597.
52. Á. Pelayo, S. Vũ Ngọc: Constructing integrable systems of semitoric type, *Acta Math.* 206 (2011) 93-125.
53. Á. Pelayo, S. Vũ Ngọc: Symplectic theory of completely integrable Hamiltonian systems, *Bull. Amer. Math. Soc. (N.S.)* 48 (2011) 409-455.
54. Á. Pelayo, S. Vũ Ngọc: First steps in symplectic and spectral theory of integrable systems, *Discrete and Cont. Dyn. Syst., Series A* **32**, 3325–3377 (2012).
55. F. Santos: Non-connected toric Hilbert schemes, *Mathematische Annalen*, 332:3 (2005), 645–665.
56. R. Schneider: *Convex Bodies: The Brunn-Minkowski Theory* (1st ed.), Encyclopedia of Mathematics and its Applications, Cambridge University Press, 1993
57. I. Schur: Über eine Klasse von Mittelbildungen mit Anwendungen auf die Determinantentheorie. *Sitzungsber. Berl. Math. Ges.* 22 (1923) 9-20.
58. G.C. Shephard, R.J. Webster: Metrics for sets of convex bodies. *Mathematika* 12, 73–88 (1965).
59. J. Włodarczyk: Decomposition of birational toric maps in blow-ups and blow-downs. A proof of the Weak Oda Conjecture. *Trans. Am. Math. Soc.* 349, 373–411 (1997)
60. G. M. Ziegler: *Lectures on Polytopes*, Springer, 1994.

ÁLVARO PELAYO, FACULTAD DE CIENCIAS MATEMÁTICAS, UNIVERSIDAD COMPLUTENSE DE MADRID, 28040 MADRID, SPAIN AND REAL ACADEMIA DE CIENCIAS EXACTAS, FÍSICAS Y NATURALES, SPAIN.
E-mail address: `alvpe101@ucm.es`

FRANCISCO SANTOS, DEPARTAMENTO DE MATEMÁTICAS, ESTADÍSTICA Y COMPUTACIÓN, UNIVERSIDAD DE CANTABRIA, AV. DE LOS CASTROS 48, 39005 SANTANDER, SPAIN
E-mail address: `francisco.santos@unican.es`