RECURSIVELY REGULAR SUBDIVISIONS AND APPLICATIONS*

Rafel Jaume† *and Günter Rote*†

Abstract. We generalize regular subdivisions (polyhedral complexes resulting from the projection of the lower faces of a polyhedron) introducing the class of *recursively regular subdivisions*. Informally, a recursively regular subdivision is a subdivision that can be obtained by splitting some faces of a regular subdivision by other regular subdivisions (and continue recursively). We also define the *finest regular coarsening* and the *regularity tree* of a polyhedral complex. We prove that recursively regular subdivisions are not necessarily connected by flips and that they are acyclic with respect to the in-front relation. We show that the finest regular coarsening of a subdivision can be efficiently computed, and that whether a subdivision is recursively regular can be efficiently decided. As an application, we also extend a theorem known since 1981 on illuminating space by cones and present connections of recursive regularity to tensegrity theory and graph embedding problems.

1 Introduction

Regular polyhedral complexes appear in a wide variety of situations. The minimization diagram of a set of linear functions, whose regularity follows almost directly from the definition, is a common instance. Power diagrams are regular complexes as well. It is not hard to see that an arrangement of hyperplanes is a regular subdivision as well; it is the projection of the lower envelope of the dual of a zonotope [17]. Yet another remarkable example is the Delaunay triangulation of a point set.

Regular subdivisions are quite well-understood even in higher dimensions. Although, as shown by Santos [29], the triangulations of a point set in dimension five and higher are not always connected via flips, regular triangulations are. Another remarkable result, which holds in any dimension, is that regular subdivisions contain no cycles in the visibility relations in the sense of Edelsbrunner [18]. On the other hand, not so much is known about non-regular subdivisions. Several generalizations of regularity have been studied in order to better understand them. For instance, the subdivisions induced by the projection of a polytope onto another polytope, introduced by Billera and Sturmfels [8], have been extensively studied together with their variants (see, for instance, Pournin [26]). In a different direction, Takeuchi [30] investigated a subclass of non-regular triangulations which is related to a strengthening of Edelsbrunner's acyclicity criterion.

*Research by R. Jaume was supported by "Obra Social la Caixa" and DAAD. Research by G. Rote was partially supported by the ESF programme EuroGIGA-VORONOI, Deutsche Forschungsgemeinschaft (DFG), project RO 2338/5-1.

†*Freie Universität Berlin, Institut für Informatik, Takustraße 9, 14195 Berlin, Germany* jaume@inf.fu-berlin.de, rote@inf.fu-berlin.de

We will use the notation $[n]$ to refer to the set $\{1, \ldots, n\}$. The d-dimensional Euclidean space will be denoted by $\mathbb{R}^d$ and $\|\cdot\|$ will denote the Euclidean norm.

1.1 Polyhedral complexes, fans and subdivisions

We assume that the reader is familiar with the notions of polytope, polyhedron, polyhedral complex, cone and fan (see [15, Chapter 2] for a detailed discussion on this topic). Nevertheless, we state some definitions and notations that differ from the most commonly used ones in the literature.

The polyhedra in a polyhedral complex will be called *faces*. The top-dimensional faces of a pure polyhedral complex are called *cells*, and the faces of one-dimension lower *walls*. The set of cells of a polyhedral complex $\mathcal{S}$ will be denoted by *cells*$(\mathcal{S})$. Given two complexes $\mathcal{S}$ and $\mathcal{S}'$, we say that $\mathcal{S}'$ *refines* $\mathcal{S}$ if every face of $\mathcal{S}'$ is contained in a face of $\mathcal{S}$. We say that $\mathcal{S}'$ is a *refinement* of $\mathcal{S}$ and $\mathcal{S}$ is a *coarsening* of $\mathcal{S}'$.

A polyhedral fan is *complete* if the union of all its cones is the whole ambient space. The (unbounded) one-dimensional faces of a polyhedral fan will be called *rays*. The *restriction* of a d-dimensional polyhedral complex $\mathcal{S}$ to a polyhedron $Q \subset \mathbb{R}^d$ is the polyhedral complex consisting of the faces $Q \cap F$ for all faces $F \in \mathcal{S}$. We let $\mathcal{S}|_Q$ denote the restriction of $\mathcal{S}$ to Q.

A *polyhedral subdivision* (or *subdivision*, for short) of a finite set of points $A \subset \mathbb{R}^d$ is a polyhedral complex whose vertices belong to A and the union of whose cells is the convex hull of A. Each cell C is then the convex hull of $C \cap A$. A subdivision of a finite set V of vectors is a polyhedral fan whose rays have as directions vectors of V and whose union is the positive span of V. A *triangulation* is a subdivision of a point set consisting only of simplices. The set of subdivisions of a point (or vector) set form a poset with the refinement relation. Our notion of a subdivision and the refinement relation are simpler than the more subtle definitions in [15, Section 2.3] or in [29], but the differences are not relevant to our results.

1.2 Regular subdivisions

Given a point $a \in \mathbb{R}^d$ and a scalar $\lambda \in \mathbb{R}$ we denote by $\binom{a}{\lambda} \in \mathbb{R}^{d+1}$ the tuple (regarded as a point) resulting from adding the coordinate λ to a.

Definition 1.1. Let $A \subset \mathbb{R}^d$ be a finite set of points. A subdivision $\mathcal{S}$ of A is *regular* if there exists a *height function* $\omega \colon A \to \mathbb{R}$ such that each face of $\mathcal{S}$ is the projection of a face in the lower convex hull of

$$A^\omega = \left\{ \begin{pmatrix} a \\ \omega(a) \end{pmatrix} : a \in A \right\}.$$

The function ω will be identified with the vector $\omega \in \mathbb{R}^A$. The notation A^ω will be used as a function of a point set A and a height function or vector ω. Given a cell

$C \in cells(\mathcal{S})$, we will also use the notation

$$A^\omega|_C = \left\{ \begin{pmatrix} a \\ \omega(a) \end{pmatrix} : a \in A \cap C \right\}.$$

The following proposition indicates that the regularity of a polyhedral subdivision can be expressed locally. We refer to [15, Chapter 5] for details.

Proposition 1.2. *(Folding form [15, Section 5.2]) Let $A \subset \mathbb{R}^d$ be a finite set of points. A polyhedral subdivision $\mathcal{S}$ of A is regular if there exists a height function $\omega : A \to \mathbb{R}$ such that for every cell $C \in cells(\mathcal{S})$, the points of $A^\omega|_C$ lie in a hyperplane* (coplanarity condition), *and for every wall $W = C \cap D$, where $C, D \in cells(\mathcal{S})$, the point $\binom{a}{\omega(a)}$ lies strictly above the hyperplane containing $A^\omega|_D$, for all $a \in A \cap (C \setminus D)$* (local folding condition).

In view of the previous result, it is easy to see that the regularity of a subdivision is equivalent to the feasibility of a linear program. We sketch the proof of this well-known fact because we will use the notation later.

Note that the coplanarity condition for a cell C can be translated into a set of linear homogeneous equations in the heights of its vertices. Indeed, it is enough to choose a spanning set of vertices $B = \{b_1, \ldots, b_{d+1}\}$ of C and require

$$\begin{vmatrix} 1 & \ldots & 1 & 1 \\ b_1 & \ldots & b_{d+1} & a \\ \omega(b_1) & \ldots & \omega(b_{d+1}) & \omega(a) \end{vmatrix} = 0, \tag{1}$$

for every vertex $a \in A \cap (C \setminus B)$.

By developing the last row of the determinant, it becomes clear that these conditions are linear equations in the heights of the lifted points. Hence, all coplanarity conditions together restrict the set of possible height functions ω to a linear subspace of $\mathbb{R}^n$.

Consider now the local folding condition for a wall $W = C \cap D$ with $C, D \in cells(\mathcal{S})$. Let $\{b_1, \ldots, b_{d+1}\}$ be a spanning set of vertices of D, and let $a \in A \cap (C \setminus D)$. The local folding condition for W can be expressed as

$$\begin{vmatrix} 1 & \ldots & 1 \\ b_1 & \ldots & b_{d+1} \end{vmatrix} \begin{vmatrix} 1 & \ldots & 1 & 1 \\ b_1 & \ldots & b_{d+1} & a \\ \omega(b_1) & \ldots & \omega(b_{d+1}) & \omega(a) \end{vmatrix} > 0. \tag{2}$$

By developing the last row of the second determinant, it becomes clear that this condition is a linear homogeneous strict inequality in the heights of the lifted points. Therefore, the local folding conditions for the walls of a subdivision $\mathcal{S}$ define together a relatively-open cone in the subspace determined by the coplanarity conditions.

Definition 1.3. The *regularity system* of a subdivision $\mathcal{S}$ is the collection of equations and inequalities resulting from its coplanarity and local folding conditions. The *weak regularity system of* $\mathcal{S}$ is the system resulting from the replacement of the strict inequality in (2) with a weak inequality. The *secondary cone of* $\mathcal{S}$ is the set of solutions of the weak regularity system.

The regularity system can be defined for coarsenings of polyhedral complexes, even if they are not polyhedral complexes (that is, if the "faces" fail to be convex or the tessellation is not face-to-face). Moreover, the definitions and statements presented here can be easily translated to the case where the initial object A is a set of vectors instead of points. In such a case, the cells of the complex are cones forming a polyhedral fan whose 1-faces are rays with directions taken from A. The local folding (and coplanarity) conditions then lose the row of ones in the determinants appearing in (1) and (2), and also one column each, since affine bases must be replaced with linear bases. We will use the term subdivision in an ambiguous manner to stress this fact and focus on point-set subdivisions in the proofs.

1.3 The secondary fan and the secondary polytope

Regular subdivisions were first studied by Gelfand, Kapranov and Zelevinsky [21], who introduced the secondary fan and the secondary polytope. These two objects encode the combinatorics of the refinement poset of the regular subdivisions of a point set. We next give the necessary definitions to state this result more precisely.

Definition 1.4. The GKZ-*vector* $\alpha(\mathcal{T})$ of a triangulation $\mathcal{T}$ of a finite point set A is the vector $\alpha(\mathcal{T}) \in \mathbb{R}^A$ whose a-th component is

$$\sum_{\substack{C \in cells(\mathcal{T}) \\ C \ni a}} \mathrm{Vol}(C).$$

The convex hull $\Sigma(A) \subset \mathbb{R}^A$ of all vectors $\alpha(\mathcal{T})$ over all triangulations $\mathcal{T}$ of A is an $(n-d-1)$-dimensional polytope called the *secondary polytope of* A.

Theorem 1.5 (Gelfand, Kapranov and Zelevinsky [21])**.** *The secondary cones of the regular triangulations of a d-dimensional point set A define an $(n-d-1)$-dimensional complete polyhedral fan (called the* secondary fan of A*). The secondary fan of A is the normal fan of* $\Sigma(A)$.

As a consequence, the vertices of $\Sigma(A)$ correspond to regular triangulations of A and the edges of $\Sigma(A)$ correspond to flips between regular triangulations. This implies that the regular triangulations of A are connected in the graph of flips. (There may be additional flips between two regular triangulations which are not represented as edges of the secondary polytope [15, p. 234].)

1.4 Edelsbrunner's acyclicity theorem

Another nice (geometrically induced) combinatorial property of regular subdivisions is the acyclicity in their *in-front* relation. We state here the definitions we need for later on.

Definition 1.6. Let v be a point in $\mathbb{R}^d$ and $S, T \subset \mathbb{R}^d$ be two disjoint convex sets. We say that S is *in front of* T with respect to v if there is an open ray ρ starting at v so that $S_0 = \rho \cap S \neq \emptyset$, $T_0 = \rho \cap T \neq \emptyset$ and every point of S_0 lies between v and any point of T_0.

This relation is called the *in-front* relation (from v). It is asymmetric because of the convexity of S and T. The relation can be defined for a direction as well, when v is considered to lie at infinity. The definition is similarly extended to polyhedral fans.

Definition 1.7. A polyhedral complex is said to be *cyclic* from a point v (possibly at infinity) if the in-front relation from v over the interior of its cells contains a cycle. The complex is called *acyclic* if it is not cyclic from any point.

Theorem 1.8 (Acyclicity Theorem [18])**.** *Regular polyhedral complexes are acyclic.*

1.5 Our contribution

This article is mainly concerned with recursively regular subdivisions. Intuitively, a polyhedral complex $\mathcal{S}$ is *recursively regular* if it is regular, or it has a regular coarsening $\mathcal{S}'$ such that for each cell $C \in cells(\mathcal{S}')$, the restriction of $\mathcal{S}$ to C is recursively regular. It is easy to see that this class of subdivisions generalizes regular subdivisions.

We give a characterization of recursively regular subdivisions, which leads to efficient algorithms for their recognition and provides meaningful structural properties. In order to do this, we introduce two constructions closely related to recursively regular subdivisions: the finest regular coarsening and the regularity tree of a polyhedral subdivision (Section 2). We provide algorithms for the construction of these objects, which have applications in different areas that we explore.

In addition, we examine some of the combinatorial properties of recursively regular subdivisions in comparison to regular subdivisions. In particular, we show that, unlike the regular subdivisions, the recursively regular subdivisions of a point set are not necessarily connected by bistellar flips. On the other hand, recursively regular subdivisions remain acyclic in the sense of Definition 1.7.

As the main application, we address the problem of finding (or deciding whether it exists) a one-to-one assignment of a set of floodlights to a set of points such that the floodlights cover the space when translated to the assigned points (Section 3). The given floodlights are assumed to be the cells of a complete polyhedral fan. We say that the fan is *universal* if the floodlights can cover the space regardless of the given point set. We prove that recursively regular fans are indeed universal, and that having a cycle in visibility is sufficient yet not necessary for a fan to be non-universal. It remains open to give a characterization of universal fans.

We also examine two related graph-theoretic problems. The first one is concerned with straight line-segment embeddings of digraphs on point sets subject to certain directional constraints on the arcs (Section 4.1). We show that a big family of digraphs (together with the direction constraints) can be embedded in any given point set, whereas some non-trivial digraphs (with constraints) cannot be embedded in some types of point sets. The second deals with rigidity of tensegrity frameworks (Section 4.2). Specifically, we show how to detect the redundant (useless, in a sense) cables from a spider web.

2 The finest regular coarsening and the regularity tree

In this section, we study the finest regular coarsening of a subdivision, which we will use afterwards to define the regularity tree. Finally, we will introduce the class of recursively regular subdivisions and analyze some of its properties.

Roughly speaking, the finest regular coarsening of a subdivision is the finest among all coarsenings of the subdivision that are regular. One should note that it is not obvious whether this object is well-defined. We show first that this is indeed the case. We do it by observing that merging two cells of a subdivision corresponds to converting a local folding condition into a coplanarity condition. Furthermore, this transformation can be done by simply replacing the strict inequality by an equation with the same coefficients. In other words, we are looking for the smallest set of inequalities we need to "relax" in order to make a given system compatible.

Checking whether a subdivision is regular leads to a system of strict inequalities of the form

$$\begin{aligned} &Mx > 0 \\ &x \in \mathbb{R}^n. \end{aligned}$$

for some matrix $M \in \mathbb{R}^{m\times n}$ with row vectors $s_1, \ldots, s_m \in \mathbb{R}^n$. If this system has no solutions, we will want to relax some of the strict inequalities to weak inequalities. In general, the problem of *relaxing the smallest number of inequalities that makes an infeasible system feasible* is difficult. However, in our case, where all constraints are homogeneous, these inequalities are easy to find: they are the *equality set* $E(M)$ of the system of inequalities $Mx \geq 0$.

Definition 2.1. Let $M \in \mathbb{R}^{m\times n}$ be a matrix with row vectors $s_1, \ldots, s_m \in \mathbb{R}^n$. The *system of* M, denoted by $S(M)$, is the system

$$S(M)\colon \begin{cases} Mx > 0 \\ x \in \mathbb{R}^n. \end{cases} \tag{3}$$

The *equality set* $E(M)$ of $S(M)$ is

$$E(M) := \{i \in [m] \mid \langle s_i, x\rangle = 0 \text{ for all } x \in P\}.$$

The *smallest (compatible) relaxation of* $S(M)$ is the system

$$\begin{aligned} &\langle s_i, x\rangle = 0, \text{ for all } i \in E(M) \\ &\langle s_i, x\rangle > 0, \text{ for all } i \in [m] \setminus E(M) \\ &x \in \mathbb{R}^n. \end{aligned}$$

Clearly, if $i \in E(M)$, then we cannot fulfill $\langle s_i, x\rangle > 0$, even if we relax all the other inequalities into weak inequalities. On the other hand, there is a solution x^* which satisfies *all* constraints $\langle s_i, x\rangle > 0$ with $i \in [m] \setminus E$: By definition, there is for each $i \in [m] \setminus E$ a solution $x^{(i)}$ of $Mx^{(i)} \geq 0$ with $\langle s_i, x^{(i)}\rangle > 0$. We then set x^* to the sum of these solutions $x^{(i)}$.

These observations directly lead to the following folklore statement, which we state for reference.

Proposition 2.2. *Let $M \in \mathbb{R}^{m \times n}$. The equality set $E(M)$ of the system $S(M)$ can be computed by solving m linear programs in n variables and m constraints.*

Proof. We simply maximize, for each $i \in [m]$, the objective function $\langle s_i, x\rangle$ over the polyhedron P. The maximum is 0 if and only if i belongs to $E(M)$. □

In practice, one can try to speed up the procedure by maximizing the sum of several terms $\langle s_i, x\rangle$: those which have not yet been excluded from $E(M)$. This gives the chance of identifying more elements at a time that cannot belong to $E(M)$, and thus it potentially reduces the number of linear programming problems. To avoid having to deal with unbounded problems, one can add some additional inhomogeneous constraints.

There is an alternative approach to compute $E(M)$ via the dual problem, which successively identifies elements that must be included in $E(M)$ [24, 25].

2.1 The finest regular coarsening of a subdivision

The algebra developed above will make it very easy to show that there exists a (well-defined) finest regular coarsening of a polyhedral subdivision. We next introduce some additional terminology concerning coarsenings.

Given two coarsenings $\mathcal{S}_1$ and $\mathcal{S}_2$ of $\mathcal{S}$, we say that $\mathcal{S}_1$ is *finer* than $\mathcal{S}_2$ if $\mathcal{S}_2$ is a coarsening of $\mathcal{S}_1$. A coarsening is *proper* if it has strictly fewer cells than the original subdivision. The *trivial coarsening* is the one that merges all cells into a single one.

Using the definitions in [15, Section 2.3], the refinement relation induces a partial order on the set subdivisions. Furthermore, the restriction of this partial order to regular subdivisions is a lattice. This lattice is isomorphic to the face lattice of the secondary polytope of the point set. However, as far as we know, not much work has been done concerning coarsenings of non-regular subdivisions. The finest regular coarsening goes in that direction, and permits to map every non-regular subdivision to a regular one which is, in a specific sense, most similar to it.

Definition 2.3. The *finest regular coarsening* of a subdivision $\mathcal{S}$ of a point set A is the subdivision obtained by the projection of the lower hull of A^{ω_0}, where ω_0 is a solution of the smallest relaxation of the regularity system of $\mathcal{S}$.

The next theorem justifies the name in this definition.

Theorem 2.4. *Let $\mathcal{S}$ be a polyhedral subdivision, and $\mathcal{S}_0$ be its finest regular coarsening. Then, $\mathcal{S}_0$ is a regular coarsening of $\mathcal{S}$, and all regular coarsenings of $\mathcal{S}$ are coarsenings of $\mathcal{S}_0$.*

Proof. Let C and D be two adjacent cells of $\mathcal{S}$ and consider the local folding condition associated to them. Observe that replacing the inequality sign in Equation (2) by an equality sign, one gets an equation equivalent to Equation (1) for a spanning vertex set of D and a vertex $a \in A \cap (C \setminus D)$. This equation, together with the coplanarity conditions for C and D, forces all vertices in $A \cap (C \cup D)$ to be lifted into a hyperplane. Thus, the new system is equivalent to the regularity system of the polyhedral complex resulting from merging the cells C and D in $\mathcal{S}$.

That is, a solution to a relaxation of the regularity system of $\mathcal{S}$ corresponds to a regular coarsening of $\mathcal{S}$. In addition, the more inequalities are strictly satisfied, the finer the coarsening is. Hence, $\mathcal{S}_0$ is a coarsening of $\mathcal{S}$, it is regular and it is the finest one satisfying these conditions. □

2.2 Relation to the secondary polytope

The finest regular coarsening of subdivision of a finite point set A can also be characterized in terms of the secondary polytope $\Sigma(A)$. Considering the definitions of subdivision and refinement used in [15, Section 2.3], the faces of $\Sigma(A)$ correspond to regular subdivisions of A, and inclusion between faces corresponds to coarsening between subdivisions. The vertices of $\Sigma(A)$ are the GKZ-vectors of the regular triangulations of A. Non-regular triangulations have GKZ-vectors that are not vertices of $\Sigma(A)$, and moreover, the function α mapping non-regular triangulations to their GKZ-vectors may not be injective [15, Section 5.2]. Nevertheless, even in the non-regular case, the normal cone of $\alpha(\mathcal{T})$ in $\Sigma(A)$ (which is then not full-dimensional) is isomorphic to the secondary cone of $\mathcal{T}$ [21]. It is then not surprising that the finest regular coarsening of a triangulation $\mathcal{T}$ corresponds to the subdivision associated to the smallest face of $\Sigma(A)$ containing $\alpha(\mathcal{T})$.

The secondary cone can be defined also for general subdivisions (not only for triangulations), see [15, Section 5.2]. This cone is contained in the linear subspace L of the height-functions space defined by the coplanarity conditions. Of course, the cone is also contained in the affine hull H of the secondary fan, which is $(n-d-1)$-dimensional. Then, a subdivision is regular if and only if its secondary cone is full-dimensional in $L \cap H$. In accordance with Definition 2.3, the finest regular coarsening of a subdivision $\mathcal{S}$ is the subdivision induced by any height function in the relative interior of the secondary cone of $\mathcal{S}$. The characterization of the finest regular coarsening in the previous paragraph is equivalent to this statement, by the duality between the secondary fan and the secondary polytope $\Sigma(A)$.

2.3 The regularity tree and recursively regular subdivisions

Roughly speaking, recursively regular subdivisions are subdivisions that can be decomposed, via a regular coarsening, into recursively regular pieces.

Definition 2.5. A polyhedral subdivision $\mathcal{S}$ is *recursively regular* if it is regular or there exists a proper, non-trivial, and regular coarsening $\mathcal{S}'$ of $\mathcal{S}$ such that $\mathcal{S}|_C$ is recursively regular for each cell $C \in \mathit{cells}(\mathcal{S}')$.

This definition can be extended to polyhedral fans. We will use the notation $\mathfrak{R}(A)$ to refer to the set of recursively regular subdivisions of a point configuration A. The class of all recursively regular subdivisions of any point set will be denoted by $\mathfrak{R}$. We will show that $\mathfrak{R}$ is larger than the class of regular subdivisions.

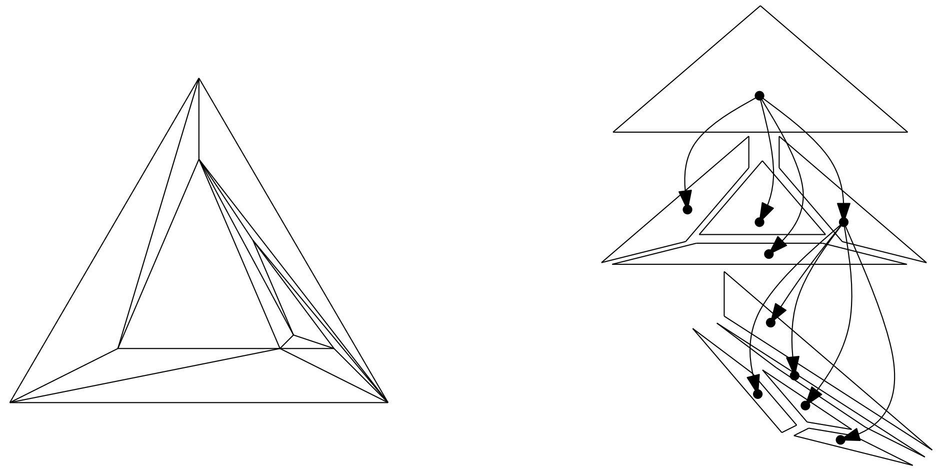

Figure 1: A recursively regular subdivision and a sketch of its regularity tree.

A specially relevant subclass of the subdivisions that do not belong to $\mathfrak{R}$ are presented next.

Definition 2.6. A subdivision is *completely non-regular* if it has more than one cell and its finest regular coarsening is its trivial coarsening.

This condition implies, in particular, that every wall of the subdivision can appear in a contradiction cycle of its regularity system.

To proceed, we need to introduce some notation and technical definitions. Given a subset $\mathcal{C}$ of $cells(\mathcal{S})$, we denote by $|\mathcal{C}|$ the ground set $\cup_{C \in \mathcal{C}} C$ covered by these cells. Similarly, if $\mathcal{S}$ is a subdivision, $|\mathcal{S}|$ will denote the union of the cells of $\mathcal{S}$.

Definition 2.7. A *subdivision tree* of a subdivision $\mathcal{S}$ of a point set A is a rooted tree such that:

i) its vertices are subsets of $cells(\mathcal{S})$

ii) its root is $cells(\mathcal{S})$

iii) if the children of $\mathcal{C}$ are $\mathcal{C}_1, \ldots, \mathcal{C}_l$, then there is a polyhedral subdivision $\mathcal{S}_{\mathcal{C}}$ of $A \cap |\mathcal{C}|$ such that $\{|\mathcal{C}_1|, \ldots, |\mathcal{C}_l|\} = cells(\mathcal{S}_{\mathcal{C}})$.

A subdivision tree is called *regular* if the subdivision $\mathcal{S}_{\mathcal{C}}$ of $A \cap |\mathcal{C}|$ used to split the node $\mathcal{C}$ is regular for each node $\mathcal{C}$ of the tree.

Note that a subdivision is recursively regular if and only if it has a regular subdivision tree. However, a subdivision can have many subdivision trees, and even many regular subdivision trees. Fortunately, we can define a canonical one, which will be later used to decide if a subdivision is recursively regular:

Definition 2.8. The *regularity tree* of the subdivision $\mathcal{S}$ is the subdivision tree created by the following recursion:

(a) If a subdivision $\mathcal{S}$ is regular or completely non-regular, its regularity tree is the tree whose single node is $|\mathcal{S}|$.

(b) Otherwise, let $\mathcal{S}_0$ be the finest regular coarsening of $\mathcal{S}$. The regularity tree is the root node $|\mathcal{S}|$ with a child for each cell $C \in \mathit{cells}(\mathcal{S}_0)$. The subtrees are regularity trees of $\mathcal{S}|_C$ for $C \in \mathit{cells}(\mathcal{S}_0)$.

Figure 1 exhibits an example of a regularity tree. The figure shows a triangulation in $\mathfrak{R}$ which needs two levels of recursion to fit the definition of recursively regular subdivision. The coordinates of this example and a proof that the finest regular coarsening of the depicted subdivision is the subdivision defined by the second level of the tree are provided in Appendix C. Note that the example consists of a "pinwheel" triangulation inserted into a triangle of a bigger copy of the pinwheel triangulation. The insertion procedure can be repeated recursively to obtain a family of triangulations where the number of levels of the regularity tree grows linearly in the number of vertices.

The leaves of the regularity tree of $\mathcal{S}$ form a partition of $\mathit{cells}(\mathcal{S})$. There are two possibilities for the subdivision in a leaf of the regularity tree: it is are either regular or completely non-regular. Accordingly, we speak of *regular leaves* and *completely non-regular leaves.*

The following theorem relates the regularity tree and the recursive regularity of a subdivision.

Theorem 2.9. *A polyhedral subdivision $\mathcal{S}$ is recursively regular if and only if the leaves of its regularity tree are regular.*

Before proceeding to the proof of the theorem, we state an easy technical lemma without proof.

Lemma 2.10. *Let $\mathcal{S}$ be a subdivision of a finite point set $A \subset \mathbb{R}^d$ and $Q \subset \mathbb{R}^d$ be a polyhedron. Then:*

i) If $\mathcal{S}$ is regular, then $\mathcal{S}|_Q$ is regular as well.

ii) If $\mathcal{S}'$ refines $\mathcal{S}$, then $\mathcal{S}'|_Q$ refines $\mathcal{S}|_Q$.

iii) If $\mathcal{S}$ is recursively regular, then $\mathcal{S}|_Q$ is recursively regular as well.

Proof of Theorem 2.9. If all leaves of the regularity tree are regular, the regularity tree itself certifies the recursive regularity of $\mathcal{S}$. That is, the hierarchy of regular coarsenings encoded by the regularity tree ensures that the subdivision satisfies the recursive conditions in Definition 2.5. The *if* direction is then proved.

For the *only if* direction, we will prove that the leaves of the regularity tree of any subdivision in $\mathfrak{R}$ are regular. We do this by induction on the number of cells of the subdivision. The base case is when the subdivision consists of a single cell C. In this case, the only leaf of its regularity tree is C, which is regular.

For the inductive step, let $\mathcal{S}$ be in $\mathfrak{R}$, and assume that the regularity tree of any subdivision in $\mathfrak{R}$ with fewer cells has regular leaves. If $\mathcal{S}$ is regular, then the statement is trivially true. Assume then that $\mathcal{S}$ is not regular and let $\bar{\mathcal{S}}$ be a non-trivial regular coarsening splitting $\mathcal{S}$ into smaller subdivisions in $\mathfrak{R}$. Such a coarsening exists by the definition of recursive regularity. Indeed, there is a regular subdivision tree of $\mathcal{S}$ representing a set of coarsenings certifying that $\mathcal{S}$ belongs to $\mathfrak{R}$. The second part of Theorem 2.4 asserts that $\bar{\mathcal{S}}$ is a coarsening of the finest regular coarsening $\mathcal{S}_0$ of $\mathcal{S}$. Then, given a cell $C \in \mathit{cells}(\mathcal{S}_0)$, there is a cell $C' \in \mathit{cells}(\bar{\mathcal{S}})$ such that $C \subset C'$. Since $\mathcal{S}|_{C'} \in \mathfrak{R}$ and $C \subset C'$, it follows from Lemma 2.10.iii that $\mathcal{S}|_C \in \mathfrak{R}$. By induction hypothesis, the leaves of the regularity tree of $\mathcal{S}|_C$ are regular. Since the leaves of the regularity tree of $\mathcal{S}$ are the leaves of the regularity trees of $\mathcal{S}|_C$ for all $C \in \mathit{cells}(\mathcal{S}_0)$, the proof is completed. □

We present now some properties of recursively regular subdivisions.

Proposition 2.11. *Let A be a finite point set.*

i) Every regular subdivision of A is recursively regular. The converse does not hold.

ii) Every recursively regular subdivision of A is acyclic. The converse does not hold.

Proof. i) Regular subdivisions are in $\mathfrak{R}$ by definition.

A non-regular triangulation that belongs to $\mathfrak{R}$ is shown in Figure 1. This example is analyzed in Appendix C. More examples may be found in [25, Figure 2a] or in [30, Figure 3] (where the recursive regularity is not analyzed).

ii) We will prove that any $\mathcal{S} \in \mathfrak{R}$ must be acyclic by induction on the number of cells. For the base case, we observe that a single-cell subdivision is always acyclic. If $\mathcal{S}$ has more than one cell, we distinguish two cases. If $\mathcal{S}$ itself is regular, then Theorem 1.8 shows that it must be acyclic. Otherwise, let $\mathcal{S}_0$ be the finest regular coarsening of $\mathcal{S}$. Assume for the sake of contradiction that $\mathcal{S}$ contains a cycle. If the cycle contains cells belonging to more than one cell of $\mathcal{S}_0$, it induces a cycle in $\mathcal{S}_0$ contradicting the fact that $\mathcal{S}_0$ is regular. If the cycle is contained a single cell $C \in \mathit{cells}(\mathcal{S}_0)$, then it is also a cycle in $\mathcal{S}|_C$. Since $\mathcal{S}|_C \in \mathfrak{R}$ and it has fewer cells than $\mathcal{S}$, this leads to a contradiction as well.

For the strictness of the inclusion, we refer to the example of Figure 2, which shows an acyclic subdivision that does not belong to $\mathfrak{R}$. These properties are established in Appendix B. □

The next proposition illustrates that $\mathfrak{R}$ includes some "pathological" triangulations. More precisely, we will show that there are triangulations in $\mathfrak{R}$ that are not connected in the graph of flips of their (common) vertex set. To prove this, we will simply show that the non-regular triangulations used by Santos [29] belong to $\mathfrak{R}$.

Proposition 2.12. *There exists a point set $A \subset \mathbb{R}^5$ whose recursively regular triangulations are not connected by geometric bistellar flips.*

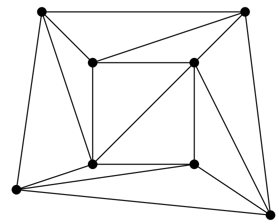

Figure 2: An acyclic triangulation that is completely non-regular.

Proof. Santos [29] constructs a set of triangulations $\mathfrak{T}$ of a five-dimensional point set A that are pairwise disconnected in the graph of flips. We show that all triangulations in $\mathfrak{T}$ are recursively regular. The convex hull of A is a prism $\mathcal{P}$ over a four-dimensional polytope Q called the 24-cell. The polytope Q has 24 facets, which are regular octahedra. All triangulations in $\mathfrak{T}$ are refinements of the prism $\mathcal{P}$ (in the sense of [15, Definition 4.2.10]) over a subdivision $\mathcal{B}$ of $A \cap Q$. The subdivision $\mathcal{B}$ is a central subdivision of Q, it is thus regular (see [15, Section 9.5]) and consists of 24 pyramids over octahedra. Therefore, the prism $\mathcal{P}$ is regular as well, because the prism over a regular subdivision is regular [15, Lemma 7.2.4]. Each triangulation in $\mathfrak{T}$ triangulates the cells of $\mathcal{B}$ in a specific way. However, since a triangulation of a pyramid is regular if and only if the triangulation induced on its base is regular [15, Observation 4.2.3], and the bases of the pyramids are regular octahedra (which are known to have only regular triangulations), the restriction of any triangulation in $\mathfrak{T}$ to any cell of $\mathcal{B}$ is regular. Hence, the restriction of any triangulation in $\mathfrak{T}$ to every cell of $\mathcal{P}$ is regular as well, since a triangulation of a prism over a simplex is regular [15, Section 6.2]. Thus, every triangulation in $\mathfrak{T}$ is recursively regular. Indeed, each triangulation in $\mathfrak{T}$ is a refinement of a regular subdivision $\mathcal{P}$, and its restriction to any cell of $\mathcal{P}$ is regular. □

In fact, the proposition shows that there is a point set A with at least 12 triangulations in $\mathfrak{R}(A)$ that are pairwise disconnected and disconnected from any regular triangulation in the graph of flips of A, as observed in [29].

2.4 Algorithms

With the help of Proposition 2.2, it becomes easy to prove that the finest regular coarsening of a subdivision can be efficiently computed. The known polynomial bound for linear programming holds only if the total number L of bits needed to encode the coefficients is counted for the input size (as in the Turing machine model).

Corollary 2.13. *Let $\mathcal{S}$ be a subdivision of a point set A in any fixed dimension and let L be the total number of bits necessary to encode the coordinates of A. The finest regular coarsening of $\mathcal{S}$ can be computed in time polynomial in L.*

Proof. It follows from the definition of the finest regular coarsening that it can be determined by finding a point ω_0 in the relative interior of the secondary cone of $\mathcal{S}$, computing the point set A^{ω_0} and its convex hull to finally project its lower faces. However, it is easier to iteratively construct it following the algorithm in Proposition 2.2 to find the smallest relaxation set of its regularity system. Whenever a constraint is relaxed (a dual variable is unrestricted), we merge the cells sharing the corresponding wall. We perform the merge operation symbolically by giving a common label to the merged cells. When the iteration ends, we construct the cells of the finest regular coarsening by computing the convex hull of the vertices of the cells with the same label. Since we assume that the dimension is constant and the vertices of the finest regular coarsening are a subset of A, the construction of the cells can be done in polynomial time.

The coefficients of the linear program come from d-dimensional determinants on the coordinates of points in A. Therefore, the number of bits needed to encode them is polynomial in L. In each iteration, a linear program is solved. The number of variables is bounded by $|A|$, and there are as many constraints as walls in $\mathcal{S}$. Therefore, the whole algorithm takes polynomial time in $|A|$ and L. $\square$

The statement in the corollary is not trivial because there are subdivisions, even in the plane, that have a linear number of simultaneous flips [20], that is, a linear number of pairs of triangular cells that can be independently merged into quadrilaterals or not. Consequently, if one would want to try out all possibilities of merging cells, these subdivisions have an exponential number of potential minimal coarsenings that one might need to test for regularity. The scenario seems even worse when it comes to recursive regularity. Fortunately, as a consequence of Theorem 2.9, this can indeed be decided in polynomial time using the procedure in Corollary 2.13.

Proposition 2.14. *Let $\mathcal{S}$ be a subdivision of a point set A in any fixed dimension and let L be the total number of bits necessary to encode the coordinates of A. Whether $\mathcal{S}$ is recursively regular can be decided in time polynomial in L.*

Proof. Theorem 2.9 ensures that we only need to compute the regularity tree of $\mathcal{S}$ to decide whether $\mathcal{S}$ belongs to $\mathfrak{R}$. This is done by computing the finest regular coarsening of subdivisions of some subsets of A. Each time we go down a level in the tree, there is one wall in the finest regular coarsening that was not in any previous finest regular coarsenings. Therefore, if we charge the computation of the finest regular coarsening to this wall, we can conclude that the number of computations is bounded by the number of walls in $\mathcal{S}$, which is polynomial if d is considered to be a constant. $\square$

3 Illumination by floodlights in high dimensions

In the last decades, a wide collection of problems have been studied concerning illumination or guarding of geometric objects. The first Art Gallery problem posed by Klee asked simply how many guards are necessary to guard a polygon. Since then, considerable research has addressed several variants of this problem, such as finding watchman routes or illuminating sets of objects. A remarkable group of problems arises when the light sources (or the surveillance devices) do not behave in the same way in all directions. In the major part of the literature, these problems are studied only in the plane. A compilation of results on this type of problem can be found in [31].

The problems we are interested in assume that each light source illuminates a polyhedral cone. The first problem we look at in this section is the *space illumination problem* in three or higher dimensions. Informally speaking, the problem asks, for a given set of floodlights and a given set of points, whether there is a placement of the floodlights on the points such that the whole space is illuminated. Afterwards, we study the generalization to higher dimensions of the *stage illumination problem*, introduced by Bose, Guibas, Lubiw, Overmars, Souvaine and Urrutia [9].

We first present a result that will be used in the subsequent proofs.

Definition 3.1. The *power diagram* of a finite set of points $Q \subset \mathbb{R}^d$ (called *sites*) with assigned weights $w\colon Q \to \mathbb{R}$ is the polyhedral complex whose cells are

$$R_q = \{x \in \mathbb{R}^d : \|x-q\|^2 - w(q) \leq \|x-q'\|^2 - w(q') \text{ for all } q' \in Q\}, \text{ for } q \in Q.$$

For every $q \in Q$, the locus R_q is a polyhedron called the *region* of q. For more details on this type of diagrams, see the survey by Aurenhammer [6]. Given a finite point set S and a set Q of weighted points, we say that an assignment $\sigma\colon S \to Q$ is *compatible with the power diagram* of Q if $s \subset R_{\sigma(s)}$, for all $s \in S$. (If a point $s \in S$ lies on the boundary between power regions, a compatible assignment can assign such a point to any region in which it lies.)

The next theorem relates power diagrams and constrained least-squares assignments.

Theorem 3.2 (Aurenhammer, Hoffmann and Aronov [4])**.** *Let Q and S be finite sets of points, and let $c\colon Q \to \mathbb{N}$ be a function such that $\sum_{q \in Q} c(q) = |S|$. Then there exist weights $w\colon Q \to \mathbb{R}$ and an assignment $\sigma\colon S \to Q$ with $|\sigma^{-1}(q)| = c(q)$ for all $q \in Q$, which is compatible with the power diagram of Q.*

If $k = |Q| \leq |S| = n$, the assignment in Theorem 3.2 can be computed in $O(k^2 n \log n)$ time by an algorithm of Alberts [2]. For the special case $|S| = |Q| = n$ and $c(q) = 1$ for all $q \in Q$, the problem becomes an assignment problem and can be thus solved using the Hungarian method in $O(n^3)$ time.

The results presented in this section use recursively regular polyhedral fans. These objects are analogous to recursively regular subdivisions of a point set with vectors instead of points as base elements. We next introduce some new definitions specific to this problem.

The *ground set* of a polyhedral complex $\mathcal{C}$, denoted by $|\mathcal{C}|$, is the union of all its cells. We say that a d-dimensional polyhedral fan is *complete* if its ground set is the whole space and that it is *conic* if the ground set is a pointed d-dimensional cone. Similarly, we will talk about the *complete case* and the *conic case* to refer to instances of the problem where the given fan is complete or conic, respectively. A wall of a fan will be called *interior* if it is not contained in the boundary of the ground set of the fan. A cone K is said to *contain a direction (or vector)* v if it contains the ray ρ_v starting at the apex of K and having direction (or direction vector) v. We will say that the direction is *interior to a cone* if ρ_v intersects the boundary of K only in its apex. We fix some vector $u \in \mathbb{R}^d \setminus \{0\}$ and declare it as the *upwards* vertical direction. For a non-vertical hyperplane $\Pi \subset \mathbb{R}^d$, we let Π^+ (respectively, Π^-) denote the closed halfspace bounded by Π and containing the direction u (respectively, $-u$). Given a set $\mathcal{H} = \{\Pi_1, \ldots, \Pi_m\}$ of hyperplanes, the *upper envelope* of $\mathcal{H}$ is defined as $\bigcap_{i\in[m]} \Pi_i^+$ and the *lower envelope* is $\bigcap_{i\in[m]} \Pi_i^-$.

For completeness, we state a well-known fact [6], which reveals the relation between power diagrams and regular polyhedral complexes.

Proposition 3.3. *The following sets are equivalent:*

i) The set of regular polyhedral complexes $\mathcal{S}$ with $|\mathcal{S}| = \mathbb{R}^d$

ii) The set of power diagrams in $\mathbb{R}^d$

iii) The set of lower envelopes of arrangements of non-vertical planes in $\mathbb{R}^{d+1}$

Definition 3.4. Let P be a polyhedron $P = (\bigcap_{i\in I} \Pi_i^+) \cap (\bigcap_{j\in J} \Pi_j^-)$, where Π_i are the hyperplanes supporting the facets of P, for $i \in I$. The *reverse polyhedron* of P, denoted by P^-, is defined as $P^- = (\bigcap_{i\in I} \Pi_i^-) \cap (\bigcap_{j\in J} \Pi_j^+)$. The *reverse fan* of a polyhedral fan $\mathcal{F}$ is the fan obtained by reversing all its faces. The *reverse cone* of a conic fan is the reverse polyhedron of its ground set.

Note that if P is a cone with apex at the origin, then $P^- = -P$, where $-P$ denotes the inversion of the set P with respect to the origin. We let $A + t$ denote the result of translating the set A by the vector t.

Definition 3.5. Given a d-dimensional complete polyhedral fan $\mathcal{F}$ with n cells and a set of n points $P \subset \mathbb{R}^d$, we say that an assignment $\sigma\colon$ *cells*$(\mathcal{F}) \to P$ is *covering* if it is one-to-one and

$$\bigcup_{C\in \mathit{cells}(\mathcal{F})} (C + \sigma(C)) \supset |\mathcal{F}|.$$

The cones of the fans represent floodlights, which have to be placed at the corresponding points. They can be translated but they are not allowed to rotate, as in other variants of the problem. Figure 3 shows an example where a covering assignment is represented by the actual translation of each cell of the fan to the corresponding point. We can see that the plane is covered (or illuminated) by the translated cells and, therefore, the assignment is covering.

We are now ready to state formally the *space illumination problem*. Given a d-dimensional polyhedral fan and a set of points in $\mathbb{R}^d$ we would like to know whether there is a covering assignment for that fan and the point set. Galperin and Galperin [19] proved that a covering assignment can be found if the fan is complete and regular, regardless of the given point set and in any dimension.

Theorem 3.6 (Galperin and Galperin [19], Rote [27])**.** *Let $\mathcal{F} \subset \mathbb{R}^d$ be a complete regular polyhedral fan consisting of n cells, and let $P \subset \mathbb{R}^d$ be a set of n points. There is a covering assignment for $\mathcal{F}$ and P.*

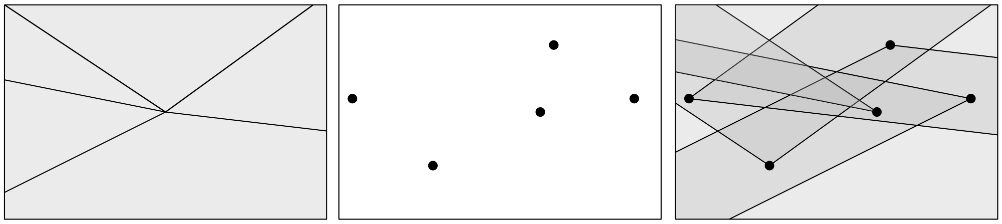

Figure 3: A covering assignment (right) for a fan $\mathcal{F}$ (left) and a point set S (center).

In particular, there is a covering assignment for a fan in the plane and any point set of the right cardinality. This last statement was rediscovered with a small variation in the formulation of the problem in [9], where an $O(n \log n)$ algorithm for finding a covering assignment is given as well. The conic case in the plane has also been considered with the extra assumption that the points are contained in the reverse cone of the fan. In this case, a covering assignment can be always found as well. However, if the points are not required to lie in the reverse cone, deciding the existence of a covering assignment becomes NP-hard even in the plane, since the problem is equivalent to the *wedge illumination problem* studied in [10].

3.1 Recursively regular fans are universally covering

We generalize first the conic case to higher dimensions and prove that it is sufficient for the fan to be recursively regular to ensure the existence of a covering assignment for any point set in the reverse cone of the fan. Afterwards, we use this result to prove that Theorem 3.6 can be extended to recursively regular fans in the complete case as well. Both generalizations are synthesized in the following statement.

Theorem 3.7. *Let $\mathcal{F} \subset \mathbb{R}^d$ be a full-dimensional recursively regular polyhedral fan consisting of n cells and $P \subset |\mathcal{F}|^-$ be a set of n points. There is a covering assignment for $\mathcal{F}$ and P.*

Note that $(\mathbb{R}^d)^- = \mathbb{R}^d$ and there is thus no restriction for P in the complete case. As a preparation for the proof, we establish two technical lemmas and a crucial lemma.

Lemma 3.8. *A conic full-dimensional fan $\mathcal{F} \subset \mathbb{R}^d$ with $|\mathcal{F}| = K$ is regular if and only if $\mathcal{F}$ is the restriction to K of a complete regular fan.*

Proof. For the *only if* direction, assume that $\mathcal{F}$ is regular and, hence, there is a cone $\widetilde{K} \subset \mathbb{R}^{d+1}$ whose lower convex hull projects on $\mathcal{F}$. This cone can be written as

$$\widetilde{K} = \left(\bigcap_{i \in I^+} \Pi_i^+ \right) \cap \left(\bigcap_{i \in I^-} \Pi_i^- \right),$$

for some hyperplanes Π_i, $i \in I^+ \cup I^-$. By convention, any vertical hyperplanes bounding $\widetilde{K}$ are considered as part of I^-. Then $\bigcap_{i\in I^+} \Pi_i^+$ is a cone whose faces project onto a complete fan $\mathcal{G}$, since the vertical direction is interior to it. By construction, $\mathcal{G}$ is regular and $\mathcal{G}|_K = \mathcal{F}$.

To prove the *if* direction, assume that $\widetilde{L} \subset \mathbb{R}^{d+1}$ is a cone whose lower hull projects onto a complete fan $\mathcal{G} \subset \mathbb{R}^d$ and let $K = \bigcap_{i\in I} \Pi_i^- \subset \mathbb{R}^d$ be a polyhedral cone. (Here, the superscript "$-$" in Π_i^- has no directional significance; Π_i^- denotes a d-dimensional half-space bounded by some hyperplane Π_i.) For every $i \in I$, let $\widetilde{\Pi}_i$ be the vertical hyperplane in $\mathbb{R}^{d+1}$ containing Π_i. Clearly the set $\widetilde{L} \cap (\bigcap_{i\in I} \widetilde{\Pi}_i^-)$ is a cone whose lower convex hull projects onto the restriction of $\mathcal{G}$ to K. □

In one step of the proof of Theorem 3.7, we want to cover a polyhedron Q by assigning n cones of a fan to n points of a point set $P \subset Q^-$. The following lemma reduces this task to finding a covering assignment for a complete fan.

Lemma 3.9. *Let $Q = \bigcap_{i\in I}(\Pi_i^+ + t_i) \subset \mathbb{R}^d$ be a full-dimensional polyhedron, where $t_i \in \mathbb{R}^d$ and $O \in \Pi_i$ for all $i \in I$. Let $\mathcal{G} \subset \mathbb{R}^d$ be a full-dimensional complete fan consisting of n cells, whose restriction $\mathcal{F}$ to the recession cone $K = \bigcap_{i\in I} \Pi_i^+$ consists of n cells as well. Let P be a set of n points with $P \subset Q^-$. If there is a covering assignment for $\mathcal{G}$ and P, then the cells of $\mathcal{F}$ translated by the corresponding assignment cover Q.*

Proof. Let $\theta\colon$ *cells*$(\mathcal{F}) \to$ *cells*$(\mathcal{G})$ be the map such that $C = \theta(C) \cap K$ for all $C \in$ *cells*$(\mathcal{F})$, and let $\sigma\colon$ *cells*$(\mathcal{G}) \to P$ be a covering assignment. We want to show that

$$\bigcup_{C\in \mathit{cells}(\mathcal{F})} (C + \sigma(\theta(C))) \supset Q. \tag{4}$$

Since σ is a covering assignment,

$$\bigcup_{D\in \mathit{cells}(\mathcal{G})} (D + \sigma(D)) = \bigcup_{C\in \mathit{cells}(\mathcal{F})} (\theta(C) + \sigma(\theta(C))) = \mathbb{R}^d \supset Q. \tag{5}$$

We will be done if we can prove that $(C + p) \cap Q = (\theta(C) + p) \cap Q$ for all $p \in Q^-$ and for all $C \in$ *cells*$(\mathcal{F})$, because then every set in the union in (5) will correspond to a set in (4). Note that $Q \subset \Pi_i^+ + p$ for any $p \in Q^-$ and for all $i \in I$, by definition of reverse polyhedron. Hence, $Q \subset K + p$. Since

$$C = \theta(C) \cap K = \theta(C) \cap \left(\bigcap_{i\in I} \Pi_i^+\right),$$

it follows that

$$(C + p) \cap Q = ((\theta(C) \cap K) + p) \cap Q = (\theta(C) + p) \cap (K + p) \cap Q = (\theta(C) + p) \cap Q.$$

Therefore, the cells of $\mathcal{F}$ translated according to $\sigma \circ \theta$ cover Q. □

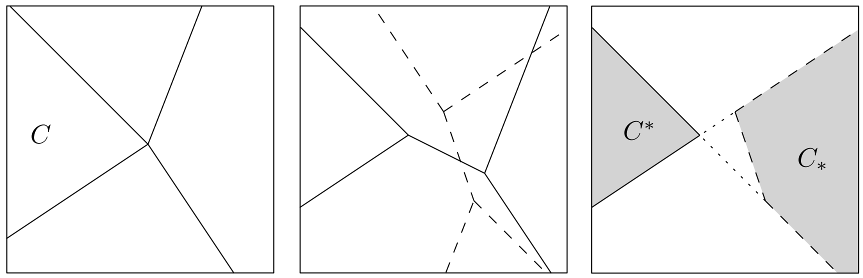

Figure 4: A fan $\mathcal{F}$ (left). An instance of $\varphi^*(\mathcal{F},\omega)$ and $\varphi_*(\mathcal{F},\omega)$ (center). A cell C_* contained in the reverse cone of C^*.

We introduce some notation and provide a proof for a generalization of a lemma from [27]. Let $\mathcal{G}$ be a regular polyhedral fan and let

$$K = \bigcap_{C \in cells(\mathcal{G})} \Pi_C^+$$

be a $(d+1)$-dimensional cone projecting onto $\mathcal{G}$, where the hyperplane Π_C supports the facet of K that projects onto C, for all $C \in cells(\mathcal{G})$. Given a function $\omega\colon cells(\mathcal{G}) \to \mathbb{R}$, let the power diagram $\varphi^*(\mathcal{G},\omega)$ be the (projection of the) upper envelope of the hyperplane arrangement obtained by vertically shifting the hyperplane Π_C by $\omega(C)$, for all $C \in cells(\mathcal{G})$. Similarly, let $\varphi_*(\mathcal{G},\omega)$ denote the projection of the lower envelope of these hyperplanes. Both power diagrams have as many cells as $\mathcal{G}$ and all of the cells are unbounded. In addition, the cells in these diagrams can be paired in a natural way with the hyperplane they come from. Let C^* denote the cell of $\varphi^*(\mathcal{G},\omega)$ corresponding to C, and let C_* denote the corresponding cell of $\varphi_*(\mathcal{G},\omega)$. These pairs of cells satisfy the following property, which is illustrated in Figure 4. We include the short proof for completeness.

Lemma 3.10 (Rote [27]). *Let $\mathcal{G}$ be a complete regular polyhedral fan, and let ω be a function $\omega\colon cells(\mathcal{G}) \to \mathbb{R}$. Then every cell $C_* \in cells(\varphi_*(\mathcal{G},\omega))$ is contained in the reverse polyhedron of $C^* \in cells(\varphi(\mathcal{G},\omega))$.*

Proof. Choose an arbitrary cell C_* of $\varphi_*(\mathcal{G},\omega)$. Let W_* be a wall between C_* and an adjacent cell D_*. W_* lies in the hyperplane h, which is the projection of $(\Pi_C+\omega(C))\cap(\Pi_D+\omega(D))$. We have to show that C_* lies on the opposite side of h as C^*. Clearly $\Pi_C+\omega(C)$ is above $\Pi_D+w(D)$ on one side of h while $\Pi_D+\omega(D)$ is above $\Pi_C+\omega(C)$ on the other side and, hence, C_* is contained in one side of h while C^* is contained in the other. $\square$

We can now prove Theorem 3.7 about recursively regular fans, which generalizes the results in [19, 27].

Proof of Theorem 3.7. We first deal with the case that $\mathcal{F}$ is a complete fan. The idea of the proof is to implicitly walk down the regularity tree and recursively split the set of cells of $\mathcal{F}$, the space in $|\mathcal{F}|$ and the points of P into smaller problems, as summarized in Algorithm 1.

Algorithm 1: Compute the covering assignment of Theorem 3.7.

Input: $\mathcal{F} \subset \mathbb{R}^d$, a full-dimensional, complete, recursively regular polyhedral fan consisting of n cells, and a set P of n points

Output: A covering assignment for $\mathcal{F}$ and P.

```
if F is regular then
|  return the covering assignment given by Theorem 3.6
else
|  F_0 ← the finest regular coarsening of F
|  foreach C ∈ cells(F_0) do
|  |  n_C ← |cells(F|_C)|
|  Find a partition of P into subsets P_C of size |P_C| = n_C, for C ∈ cells(F_0),
|      and a weight function ω such that each set P_C is contained in the
|      corresponding cell C_* of φ_*(F_0, ω), with the help of Theorem 3.2
|  foreach C ∈ cells(F_0) do
|  |  Construct a complete fan G with |cells(G)| = |cells(F|_C)| and G|_C = F|_C
|  |      by Lemma 3.8
|  |  Apply Algorithm 1 recursively to (G,P_C), yielding σ_C
|  return the combination of σ_C for all C ∈ cells(F_0)
```

If $\mathcal{F}$ is regular, Theorem 3.6 provides a covering assignment. If $\mathcal{F}$ is not regular but recursively regular, we compute the finest regular coarsening $\mathcal{F}_0$ of $\mathcal{F}$. Note that the lower envelope $\varphi_*(\mathcal{F}_0, \omega)$ is also a power diagram for some set of sites Q by Proposition 3.3. Changes in the weight function $w\colon Q \to \mathbb{R}$ of the power diagram correspond to changes in the function $\omega\colon$ *cells*$(\mathcal{G}) \to \mathbb{R}$. Therefore, we can apply Theorem 3.2 to find ω and a partition of P into one set P_C for each cell $C \in$ *cells*$(\mathcal{F}_0)$, such that P_C is contained in the cell C_* of $\varphi_*(\mathcal{F}_0, \omega)$ and $|P_C| = n_C = |$*cells*$(\mathcal{F}|_C)|$. We then want to cover recursively each cell C^* of $\varphi^*(\mathcal{F}_0, \omega)$ with the floodlights of $\mathcal{F}|_C$ and the points of P_C. Since the polyhedral cells C^* partition space, this will be sufficient.

We recursively compute a covering assignment σ_C for the complete fan $\mathcal{G}$ that extends $\mathcal{F}|_C$. Lemma 3.10 ensures that the points $P_C \subset C_*$ are in the reverse polyhedron of C^*, and therefore the hypotheses of Lemma 3.9 are satisfied. Thus, when we use σ_C as an assignment for $\mathcal{F}|_C$ and P_C, the region C^* will be covered.

The algorithm terminates because the number of cells of the fan strictly decreases in each recursive call.

In the case when $\mathcal{F}$ is a conic fan, we first construct a complete fan $\mathcal{G}$, invoking Lemma 3.8, and apply the algorithm to $\mathcal{G}$. Since $P \subset |\mathcal{F}|^-$ by assumption, we can argue as above with Lemma 3.9 that the resulting assignment covers $|\mathcal{F}|$. □

3.2 Cyclic fans are not universally covering

Let $\mathcal{F} \subset \mathbb{R}^d$ be a full-dimensional polyhedral fan with n cells. We say that $\mathcal{F}$ is *universally covering* if for any point set $P \subset \mathbb{R}^d$ of n points there exists a covering assignment for $\mathcal{F}$ and P. After showing that all recursively regular fans are universally covering, one could imagine that all fans are so. We prove that this is not the case in dimension three and higher by showing that if a fan is cyclic in some direction, there is a point set for which there is no covering assignment. This statement will easily follow from Theorem 3.15. Before proving this theorem, we need to introduce a definition and state a technical lemma.

Definition 3.11. Let $\mathcal{F} \subset \mathbb{R}^d$ be a full-dimensional polyhedral fan, let $W \in \mathcal{F}$ be common wall between two cells $C, D \in \mathit{cells}(\mathcal{F})$, and let v be a vector normal to W pointing from C to D. We say that an assignment $\sigma\colon \mathit{cells}(\mathcal{F}) \to \mathbb{R}^d$ satisfies the *overlapping condition* for W if $\langle(\sigma(C) - \sigma(D)), v\rangle \geq 0$.

An assignment satisfies this condition for a wall if and only if the copies of the two cells sharing the wall translated to the assigned points have non-empty intersection. The following observation is straightforward.

Lemma 3.12. *Let $K \subset \mathbb{R}^d$ be a full-dimensional polyhedral cone.*

(i) Any line with direction interior to K has unbounded intersection with K.

(ii) Any line with direction not contained in K has bounded intersection with K.

The next lemma follows easily.

Lemma 3.13. *Let $\mathcal{F} \subset \mathbb{R}^d$ be a full-dimensional polyhedral fan. A covering assignment for $\mathcal{F}$ must satisfy the overlapping condition for every interior wall of the fan.*

Proof. If the condition is not satisfied for the wall $H = C \cap D$, we consider a ray ρ in a direction interior to H and place it at the point $(\sigma(C) + \sigma(D))/2$. Since only the cells C and D contain the direction of ρ, only C and D can cover an unbounded part of ρ, by Lemma 3.12. In addition, none of these two cells intersect ρ. Therefore, since ρ is unbounded and we have finitely many cones, ρ cannot be completely covered. If the fan is complete, the proof is finished. Otherwise, we should note that ρ will eventually enter $|\mathcal{F}|$, since the direction of ρ is interior to an interior wall of $\mathcal{F}$ and, hence, interior to $|\mathcal{F}|$. □

Satisfying the overlapping conditions is not sufficient in general, not even in the plane. An exception is the case where all points lie on a line, which is studied in the following lemma.

Lemma 3.14. *Let $\sigma\colon \mathit{cells}(\mathcal{F}) \to P$ be an assignment for a full-dimensional polyhedral fan $\mathcal{F} \subset \mathbb{R}^d$ consisting of n cells, and a set $P \subset |\mathcal{F}|^-$ of n points lying on a line ℓ. If σ satisfies the overlapping condition for all interior walls, then it is a covering assignment.*

Proof. We prove first the complete case. We can assume without loss of generality that ℓ goes through the origin, and $\mathcal{F}$ has its apex at the origin. Fix a direction vector v parallel to ℓ. Consider any oriented line ℓ' with direction v. At infinity in direction v, ℓ' is covered by some (untranslated) cell C of $\mathcal{F}$. Hence, when C is translated to its assigned point $\sigma(C)$ of P, it still covers ℓ' at infinity because the translation is only parallel to ℓ'. Let $q \in \ell'$ be the point where ℓ' leaves $C + \sigma(C)$. Suppose that q is in the relative interior of a translated wall of C, and let $W = C \cap D$ be this wall, where $C, D \in \mathit{cells}(\mathcal{F})$. Then the overlapping condition for W and the special position of P ensures that ℓ' enters $D+\sigma(D)$ before leaving $C+\sigma(C)$. Iterating this argument, we eventually reach a cell containing the direction $-v$ that covers the unbounded remainder of ℓ'. Thus, any line ℓ' with direction v and that intersects only d- and $(d-1)$-dimensional faces of the translated cells is completely covered. The union $\mathcal{U}$ of the remaining lines with direction v (that is, the lines intersecting some $(d-2)$-dimensional face of some translated cell) is a nowhere-dense set and thus is covered as well. Indeed, for every line $\hat{\ell} \in \mathcal{U}$ we can find a line not in $\mathcal{U}$ with direction v (and, hence, covered) arbitrarily close to $\hat{\ell}$. Since the cells are closed sets, the limit of a sequence of covered lines must be covered as well, and thus $\mathcal{U}$ is covered. Since any line with direction v is covered, $\mathbb{R}^d$ is completely covered.

Assume now that $\mathcal{F}$ is a conic fan with $K = |\mathcal{F}|$. Consider a line ℓ' with direction v that enters K through a wall. Let $C \in \mathit{cells}(\mathcal{F})$ be the cell containing this wall. Since $P \subset \ell \cap K^-$, the line ℓ' enters the translated cell $C+\sigma(C)$ before entering K. The arguments for the complete case carry over unless the line crosses a wall $W+\sigma(D)$ of some translated cell $D+\sigma(D)$ such that $W \subset \partial K$. Then, again the fact that $P \subset \ell \cap K^-$ implies that the line ℓ has left K before this happens. Therefore, if ℓ' is a line with direction v that avoids $(d-2)$-dimensional faces of the translated cells (and of K), then $\ell' \cap K$ is covered. A limit argument as in the complete case ensures that then all lines with direction v have the portion intersecting K covered, and thus K is covered. □

We will now show that there is a fan and a point set for which there is no covering assignment.

Theorem 3.15. *Given a full-dimensional polyhedral fan $\mathcal{F} \subset \mathbb{R}^d$ with n cells and set of n points $P \subset |\mathcal{F}|^-$ lying on a line ℓ, there is a covering assignment for $\mathcal{F}$ and P if and only if $\mathcal{F}$ is acyclic in the direction of the line ℓ.*

Proof. Let $v \neq 0$ be a vector parallel to ℓ. We can construct a directed graph G having the cells of $\mathcal{F}$ as vertices and an edge from D to C if the vector u normal to $W = C \cap D$ pointing from C to D satisfies $\langle u, v\rangle \geq 0$.

If $\mathcal{F}$ is acyclic in the direction v, G is acyclic as well. If the points $\sigma(C)$ appear on ℓ in an order which is compatible with the partial order represented by G, then the overlapping conditions hold for σ. Lemma 3.14 ensures then that the assignment is covering.

On the other hand, if there is a visibility cycle $\tau = (C_1 \dots C_k)$ in the direction v, that is, C_i is in front of C_{i+1}, for all $i \in [k-1]$, and C_k is in front of C_1, there is a cycle in the order in which the points $\sigma(C_1), \dots, \sigma(C_k)$ should appear in the line in order to satisfy the overlapping conditions, preventing them to be satisfied for all walls of the fan. These conditions have been proved to be necessary for a covering assignment. □

If a covering assignment exists for a given point set in a line and a given fan, it can be computed in $O(n^2)$ time. We first perform a topological sort on the graph G described in the proof of Theorem 3.15. Since the number of walls is bounded by n^2, so is the number of edges, and the algorithm runs in $O(n^2)$ time. Afterwards, it only remains to sort the points

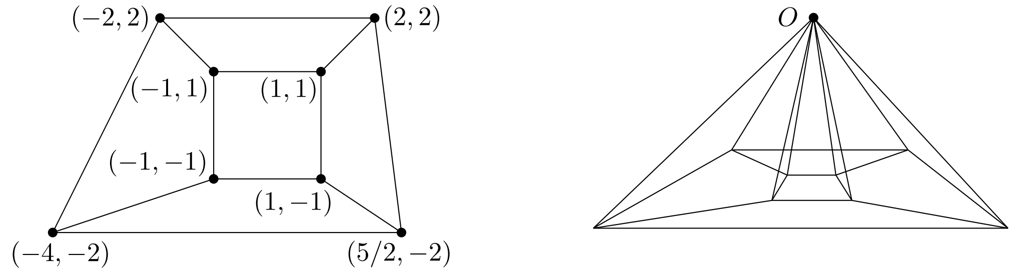

Figure 5: Acyclic subdivision that is not recursively regular (left) and the corresponding fan (right).

in P, which can be done in $O(n \log n)$ time, and assign them according to the topological sorting of G. Moreover, the topological sort algorithm would detect if the graph has a cycle and, therefore, if there is no covering assignment.

In view of the previous theorem, one might be tempted to conjecture that being acyclic is equivalent to being universally covering. We exhibit next an example to show that this is not the case.

Proposition 3.16. *There is an acyclic full-dimensional polyhedral fan $\mathcal{F} \subset \mathbb{R}^3$ consisting of n cells, and a set of n points $P \subset |\mathcal{F}|^-$ for which there is no assignment satisfying the overlapping conditions. Hence, by Lemma 3.13, there is no covering assignment for $\mathcal{F}$ and P.*

Proof. We will provide a three-dimensional fan $\mathcal{F}$ with five cells and a point set $P \subset \mathbb{R}^3$ for which there is no covering assignment. More precisely, it can be shown that for each of the 5! possible assignments, one of the eight overlapping conditions is violated.

To construct $\mathcal{F}$, take the subdivision sketched in Figure 5 (left) and embed it in the plane $\{(x, y, z) \in \mathbb{R}^3 : z = -1/8\}$. Take then the cones from the origin to each of the cells of this subdivision forming the fan displayed in Figure 5 (right).

Let P be the point set consisting of the points

$$p_1 = (29, 95, 89),\ p_2 = (55, 19, 92),\ p_3 = (54, 10, 82)$$
$$p_4 = (78, 2, 68),\ p_5 = (15, 40, 92).$$

There is no assignment for this point set fulfilling all overlapping conditions. The computations supporting this claim are given in [25, 24]. We prove in Appendix A that $\mathcal{F}$ is acyclic. □

The fan in this proof was selected because it is the fan with the fewest cells that is not recursively regular we are aware of. The point set P was then found with the help of a computer. We generated many pseudo-random samples of five points in $\mathbb{R}^3$, trying different ranges for the coordinates and several parameters for the distribution.

This last example motivates the conjecture that a fan is covering if and only if it is recursively regular. Note that a fan that is not recursively regular must have a completely non-regular convex region, and this fact could perhaps be used to construct a point set for which no covering assignment exists.

Illuminating a stage. The problem of illuminating a pointed cone using floodlights is closely related to the problem of illuminating a stage considered in [9, 13, 16, 23]. Informally, the problem in the plane asks whether given n angles and n points, floodlights having the required angles can be placed on the points in a way that a given segment (the stage) is completely illuminated. The problem can be generalized to higher dimensions where our results on covering a cone by a conic fan have new implications [24, 25].

4 Other applications and related problems

In this section we describe applications of the theoretical results introduced before.

4.1 Embeddings of directional graphs

As shown in Section 3, for the existence of a covering assignment it is necessary that there is an assignment satisfying the overlapping condition for every interior wall of the fan. Moreover, the examples we have found so far of polyhedral fans and point sets for which there is no covering assignment fail to fulfill the overlapping conditions. Hence, it could be that these conditions are also sufficient. In any case, we think that it is of independent interest to study these conditions alone, which are connected to a problem on graph embedding.

Note first that the overlapping condition for a wall can be expressed as a requirement on the order in which the two involved points are swept by a hyperplane parallel to the wall. The problem we study here asks whether, given a set of relations of this type (stated on labels) and a point set, we can find a one-to-one labeling of the point set such that every relation is satisfied. We next describe the problem formally.

Definition 4.1. A *d-dimensional directional graph* is a tuple $\vec{G} = (V, h)$, where V is a set and $h\colon V \times V \to \mathbb{R}^d$ is a function such that $h(v, u) = -h(u, v)$, for all $v, u \in V$.

The elements of V are called *vertices*. We say that $u, v \in V$ are connected by an *edge* if $h(u, v) \neq 0$. We may regard this structure as a directed graph with a non-zero direction associated to every edge. Such a graph will be called the *underlying graph* of the directional graph. Note that the condition in the definition already implies that $h(v, v) = 0$, for all $v \in V$.

Definition 4.2. 1. An *embedding* of a d-dimensional directional graph $\vec{G} = (V, h)$ on a point set $P \subset \mathbb{R}^d$ is a one-to-one assignment $\sigma\colon V \to P$ such that

$$\langle h(u, v), \sigma(v) - \sigma(u)\rangle \geq 0, \text{ for all } u, v \in V.$$

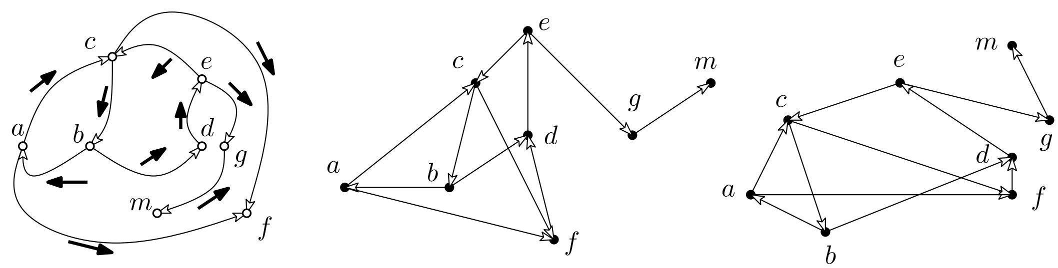

Figure 6: A directional graph (left), a drawing (center), and an embedding (right).

If such an embedding exists, we say that $\vec{G}$ is *embeddable* in P.

$\vec{G}$ is *universally embeddable* if it is embeddable on any point set $P \subset \mathbb{R}^d$ with $|P| = |V|$.

2. A *drawing* of a directional graph $\vec{G} = (V, h)$ is a bijection $\pi\colon V \to S \subset \mathbb{R}^d$ such that for all $u, v \in V$ with $h(u, v) \neq 0$, there is some $\lambda_{uv} > 0$ such that $\pi(v) - \pi(u) = \lambda_{uv} \cdot h(u, v)$.

   $\vec{G}$ is *drawable* if it has a drawing.

A directional graph is shown in Figure 6, together with a drawing and an embedding. The arrows near the edges indicate the directions associated with them. Observe that the embedding condition for an edge restricts its direction to a halfspace, while the drawing condition fixes its direction completely. Note also that the lengths of the vectors assigned by h are irrelevant for the existence of an embedding or a drawing of a directional graph. Therefore, we will consider two directional graphs (V, h) and (V, h') *equivalent* if $h(u, v)$ is a positive scalar multiple of $h'(u, v)$ for all $u, v \in V$.

Definition 4.3. 1. The *directional graph of a polytope* is the set of its vertices, together with the function $h(u, v) = v - u$ if u and v are endpoints of an edge of the polytope, and $h(u, v) = 0$ otherwise.

2. The *normal graph* of a polyhedral fan is the set of its cells with the function $h(C, D)$ being a vector normal to the wall common to C and D and pointing "from C to D" if they share a wall, and $h(u, v) = 0$ otherwise.

Note that the directional graph of a polytope and the normal graph of its normal fan are equivalent. This is a consequence of the duality between a polytope at its normal fan. The following proposition shows that these graphs provide a large family of universally embeddable directional graphs.

Proposition 4.4. *1. If a directional graph is drawable, then it is universally embeddable. In particular, a directional graph $\vec{G} = (V, h)$ whose underlying graph G is a tree is universally embeddable, regardless of h.*

2. The directional graph of a polytope is universally embeddable.

Proof. Given a point set P with $|P| = |V|$, consider a drawing π of $\vec{G}$, and let μ be the least-squares optimal matching between $\pi(V)$ and P. We will show that $\mu \circ \pi$ is an embedding of $\vec{G}$. Assume that this is not the case. Then there must be a pair $u, v \in V$ such that $\langle h(u,v), \mu(\pi(v)) - \mu(\pi(u))\rangle < 0$. Since $\pi(v) - \pi(u) = \lambda_{uv} \cdot h(u,v)$ for some $\lambda_{uv} > 0$, it follows that $\langle \pi(v) - \pi(u), \mu(\pi(v)) - \mu(\pi(u))\rangle < 0$, which contradicts the optimality of μ because swapping the images of $\pi(u)$ and $\pi(v)$ would improve the matching.

Directional graphs having a tree as underlying graph are trivially drawable and directional graphs of polytopes have the 1-skeleton of the polytope as a drawing. □

It is not hard to see that if there is a sequence of vertices $v_1, \dots, v_l, v_{l+1} = v_1$ in V and a vector $\delta \in \mathbb{R}^d$ such that $\langle h(v_i, v_{i+1}), \delta\rangle > 0$, for all $i \in [l]$, then the graph is not drawable. Such a cycle is called a *(δ-)forcing cycle*. However, the converse is not true in general: for instance, the normal graph of the subdivision in Figure 5 has no forcing cycle but it is also non-drawable.

The following proposition summarizes several relations between recursive regularity, drawability and embeddability of directional graphs. The *projection* of a d-dimensional directional graph $\vec{G}$ into a k-dimensional linear subspace $L \subset \mathbb{R}^d$ is the k-dimensional directional graph obtained by projecting each vector $h(u,v) \in \mathbb{R}^d$ onto $L \cong \mathbb{R}^k$.

Proposition 4.5.

- *(i) A projection of a universally embeddable directional graph is universally embeddable.*
- *(ii) Normal graphs of recursively regular fans are universally embeddable.*
- *(iii) Universally-embeddable graphs are not necessarily drawable.*
- *(iv) Graphs with forcing cycles are not universally embeddable.*
- *(v) There are graphs with no forcing cycles that are not universally embeddable.*

Proof.

(i) Let $\vec{G} = (V, h)$ be a d-dimensional universally-embeddable directional graph, and let L be a k-dimensional linear subspace of $\mathbb{R}^d$ with a basis $\{l_1, \dots, l_k\}$. Let $\bar{G} = (V, \bar{h})$ be the projection of $\vec{G}$ onto L, which is identified with $\mathbb{R}^k$ through the bijection

$$\begin{aligned} i\colon \mathbb{R}^k &\to L \subset \mathbb{R}^d \\ (x_1, \dots, x_k) &\longmapsto \sum_{j \in [k]} x_j l_j. \end{aligned}$$

Consider any set of $|V|$ points $\bar{P} \subset \mathbb{R}^k$, and the associated point set $P = i(\bar{P}) \subset \mathbb{R}^d$. If $\sigma\colon V \to P$ is an embedding of $\vec{G}$ on P, then $\bar{\sigma} = i^{-1} \circ \sigma$ is an embedding of $\bar{G}$ on $\bar{P}$. Indeed, $\langle h(u,v), \sigma(v) - \sigma(u)\rangle = \left\langle \bar{h}(u,v), \bar{\sigma}(v) - \bar{\sigma}(u)\right\rangle$ for all $u, v \in V$, because $\sigma(u) - \sigma(v) \in L$ and thus only the projection of $h(u,v)$ onto L contributes to the scalar product.

(ii) Let $\mathcal{F} \subset \mathbb{R}^d$ be a full-dimensional polyhedral fan consisting of n cells. Theorem 3.7 ensures that there is a covering assignment for $\mathcal{F}$ and any set P of n points. This assignment must satisfy the overlapping condition for each wall of the fan, which is equivalent to the embedding condition for the corresponding edge in its normal graph.

(iii) The normal graph of a fan is drawable if and only if the fan is regular (see, for instance, [5]). Thus, the normal graph of a recursively regular non-regular fan is not drawable. It is, however, universally embeddable, as shown in (ii).

(iv) Consider a δ-forcing cycle $v_1, \ldots, v_l, v_{l+1} = v_1$. Take a set of different points in a line having direction vector δ and label them increasingly with respect to their scalar products with δ. For any embedding σ, $\sigma(v_{i+1})$ must have a label larger than $\sigma(v_i)$, for all $i \in [l]$, which is obviously impossible.

(v) The normal graph of the fan obtained by taking cones from the subdivision in Figure 5 has no forcing cycle, since it is acyclic (in the sense of Definition 1.7). However, we have given a set of points for which all assignments violate an overlapping condition. Hence, there is no embedding of its normal graph on this point set. □

4.2 Redundancy in spider webs

We present now a problem in tensegrity theory related to the finest regular coarsening of subdivisions in $\mathbb{R}^2$. We first review the main results we will need.

Tensegrity theory studies the rigidity properties of frameworks made of bars, cables and struts. An *abstract framework* $G = (V; B, C, S)$ is a graph on the vertex set $V = \{v_1, \ldots, v_n\}$ whose edge set E is partitioned into sets B, C and S. The edges in B are called *bars*, the ones in C are called *cables* and the ones in S are called *struts*. They represent links supporting any stress, non-negative stresses, and non-positive stresses, respectively.

Definition 4.6. A (*tensegrity*) *framework* in $\mathbb{R}^2$ is an abstract framework together with an embedding of the vertices $p\colon V \to \mathbb{R}^2$.

The framework associated to an abstract framework $G = (V; B, C, S)$ will be denoted by $G(p)$ and p will be thought of as a point $(p_1, \ldots, p_n) \in \mathbb{R}^{2n}$, with $p_i = p(v_i)$ for $i \in [n]$. The configuration space $X(p)$ of $G(p)$ is

$$\begin{aligned} X(p) = \{(x_1, \ldots, x_n) \in \mathbb{R}^{2n} : \|x_i - x_j\| &= \|p_i - p_j\|, \text{ for all } v_iv_j \in B; \\ \|x_i - x_j\| &\leq \|p_i - p_j\|, \text{ for all } v_iv_j \in C; \\ \|x_i - x_j\| &\geq \|p_i - p_j\|, \text{ for all } v_iv_j \in S\}. \end{aligned} \tag{6}$$

That is, $X(p)$ is the set of embeddings of G preserving the length of the bars, making the lengths of the cables no longer and the lengths of the struts no shorter than their lengths induced by p.

A tensegrity framework $G(p)$ is *rigid* in $\mathbb{R}^d$ if there exists an open neighborhood $U \subset \mathbb{R}^{2n}$ of p such that $X(p) \cap U = M(p) \cap U$, where

$$M(p) = \{(x_1, \ldots, x_n) \in \mathbb{R}^{2n} : \|x_i - x_j\| = \|p_i - p_j\|, \text{ for all } i, j \in [n]\}$$

is the manifold of rigid motions associated to p. In other words, a framework is rigid if its only motions respecting the constraints (6) are the motions that rigidly move the whole framework. The study of the quadratic constraints in the definition of $X(p)$ can be complicated. Because of this, the notion of infinitesimal rigidity was introduced, which captures the rigidity constraints up to the first order. Consider the system of linear equations and inequalities obtained by differentiating the constraints in (6). If the solutions of the system correspond only to differentials of Euclidean motions, the framework is *infinitesimally rigid*. It is known that infinitesimal rigidity implies rigidity, and the converse is not true.

Definition 4.7. Given a framework $G(p)$, we say that $\omega\colon E \to \mathbb{R}$ is a *proper* (*equilibrium*) *stress* for $G(p)$ if the following conditions hold:

(1) $\omega(v_iv_j) = 0$ if $v_iv_j \notin E$.

(2) $\omega(v_iv_j) \geq 0$ if $v_iv_j \in C$.

(3) $\omega(v_iv_j) \leq 0$ if $v_iv_j \in S$.

(4) Every $v_i \in V$ is *in equilibrium*, that is, $\sum_{v_j \in V} \omega(v_iv_j)(p_j - p_i) = 0$.

We say that ω is *strictly proper* if the stresses on all cables and struts are non-zero.

Intuitively, ω is a proper equilibrium stress for $G(p)$ if the forces (represented by ω) exerted by the incident edges on each vertices add up to zero, taking into account that cables can support only non-negative stresses and struts can support only non-positive ones. Clearly, the stress assigning zero to all edges is proper. This stress is called the *trivial stress*.

We can associate a framework $G(\mathcal{S})$ with a given subdivision $\mathcal{S}$ in the plane by making the vertices and edges of $\mathcal{S}$ to induce the vertices and edges of the underlying graph of $G(\mathcal{S})$ and setting p to place each vertex of $G(\mathcal{S})$ on the corresponding vertex of $\mathcal{S}$. Note that it also should be prescribed which edges of $\mathcal{S}$ should be regarded as cables, bars and struts. We restrict the study to this type of frameworks for simplicity, although the results admit a more general setting (allowing non-convex faces, for instance).

We state next a special case of the Maxwell-Cremona correspondence, which relates liftings from a subdivision to a polyhedral terrain and stresses of the associated framework. When two adjacent faces of the subdivision are lifted to a terrain, the vertices of one lifted face that do not belong to the other one may be below, above or on the plane containing the other lifted face. These three types of edges are called *valleys*, *mountains* and *flat edges*, respectively. We refer to [12] for more details.

Theorem 4.8 (Maxwell-Cremona correspondence)**.** *Let $\mathcal{S}$ be a subdivision in the plane and let $G(\mathcal{S})$ be the associated framework. There is a bijection between proper stresses for $G(\mathcal{S})$ and polyhedral terrains (with one arbitrarily chosen but fixed face at height zero) projecting on $\mathcal{S}$. Positive stress values correspond to valleys, negative stress values correspond to mountains and zero stress values correspond to flat edges in the terrain.*

Stresses on a framework are also related to its (infinitesimal) rigidity.

Lemma 4.9 (Roth and Whiteley [28]). *If a tensegrity framework is infinitesimally rigid, then it has a strictly proper stress.*

We are interested in a kind of frameworks that are called *spider webs*.

Definition 4.10. A *spider web* is a framework $G(p)$ (in $\mathbb{R}^2$) consisting only of cables and whose graph is connected.

It is considered that the vertices in the convex hull of $p(V)$ are pinned down (that is, they are in equilibrium by definition).

Lemma 4.11 (Connelly [11]). *If a spider web has a strictly proper stress, then it is rigid.*

Note that the Maxwell-Cremona correspondence implies that a spider web $G(\mathcal{S})$ has a strictly-proper stress if and only if $\mathcal{S}$ is regular.

Proposition 4.12. *Let $\mathcal{S}$ be a subdivision associated and $G(\mathcal{S})$ be the associated spiderweb.*

(i) Only the cables of G corresponding to edges of the finest regular coarsening of $\mathcal{S}$ support a positive stress in any equilibrium stress of $G(\mathcal{S})$.

(ii) If $\mathcal{S}$ is recursively regular, then $G(\mathcal{S})$ is rigid.

Proof.

(i) Since the vertices on the boundary of $|\mathcal{S}|$ are in equilibrium by definition, we can make edges on this boundary to support a positive stress. For the remaining edges, Theorem 2.4 implies that if they are omitted in the finest regular coarsening, then they are lifted to flat edges by any convex lifting. Then, the Maxwell-Cremona correspondence indicates that the corresponding cables will receive zero stress in any proper equilibrium.

(ii) The finest regular coarsening $\mathcal{S}_0$ of $\mathcal{S}$ corresponds to a set of cables such that there is an equilibrium stress assigning positive values to all of them. Therefore, the spider web defined by this set of cables is rigid by Lemma 4.11. The recursively regular subdivision $\mathcal{S}|_C$ for an arbitrary cell $C \in \mathit{cells}(\mathcal{S}_0)$ can be considered an spiderweb on its own, since the vertices in the boundary of $|\mathcal{S}|_C|$ are now fixed. Thus, the argument applies recursively until a regular subdivision is reached, which are rigid by Lemma 4.11. □

Figure 7 illustrates the result. The spider web represented in it is constructed from a triangulation appearing in [1]. The edges omitted in the picture to the right, which do not belong to the finest regular coarsening of the associated subdivision, support no stress in any equilibrium. Therefore, they can be considered redundant.

Even though recursively regular subdivisions are associated to rigid spider webs, they might be far from infinitesimally rigid. For instance, if a regular subdivision is refined

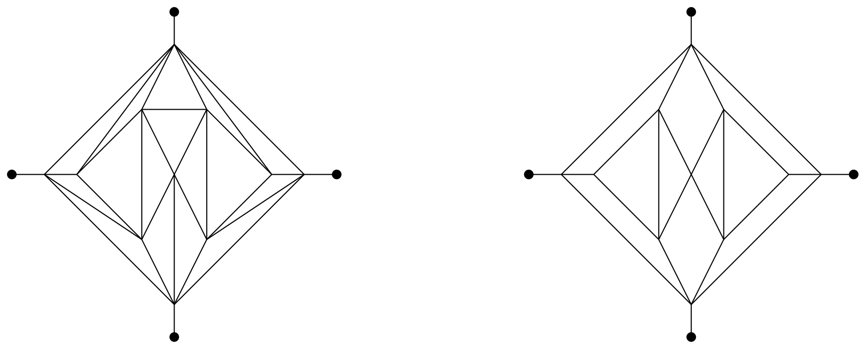

Figure 7: A spider web with 4 pinned vertices (left) and the result of removing redundant cables (right).

by adding an edge whose endpoints are interior to previous edges, the result is recursively regular but obviously not infinitesimally rigid. We next translate a well-known fact of infinitesimal rigidity to the language of finest regular coarsenings.

Corollary 4.13. *Let $\mathcal{S}$ be a subdivision in the plane and $G(\mathcal{S})$ be the associated spiderweb. If $G(\mathcal{S})$ is infinitesimally rigid, then $\mathcal{S}$ is its own finest regular coarsening.*

Proof. As Lemma 4.9 states, if a framework is infinitesimally rigid, it has a strictly-proper stress. The edges omitted in the finest regular coarsening of the associated subdivision cannot participate in such stress. Therefore, none of the edges are omitted in the finest regular coarsening of the subdivision. □

5 Concluding remarks and open problems

We have shown that the finest regular coarsening of a subdivision, which can be seen as the regular subdivision that is closest to it, can be used to define a structure called the regularity tree. The leaves of this tree define a partition of the subdivision in sub-subdivisions that are either regular or completely non-regular. The regularity tree reflects thus some of the structure of non-regular subdivisions, and it measures, in a sense, the degree of regularity. As a consequence, the class of recursively regular subdivisions arises in a natural way. We have shown that this class goes beyond regular subdivisions while excluding cyclic ones. However, we have proven that they are in general not connected by flips.

In addition, we have studied a collection of related applications, and we expect to find even more, since any theorem or algorithm based on the regularity of a subdivision and admitting a recursive scheme might be extended to apply for the larger set of recursively regular subdivisions.

In particular, we have focused on the problem of illuminating space by floodlights. It was known that regular fans are universal and our aim was to answer the question for the other fans. We have proved that not only regular fans are universal and that not only cyclic ones are non-universal. It makes then sense to ask what is the complexity class of the general problem of deciding whether the space can be covered by a given fan from a given point set (in dimensions bigger than two). It remains open as well to explore the precise limits of universality, that is, to characterize the polyhedral fans that can cover the space from any point set. A reasonable candidate is recursive-regularity. Indeed, the fact that a subdivision that is not recursively regular must have a convex sub-subdivision which is completely non-regular could be the first step towards a proof for this fact. Our results on covering the space by floodlights have implications for a three-dimensional version of the *stage illumination problem*. In data visualization, recursive partitions using regular subdivisions (Voronoi treemaps) have been used by Balzer and Deussen [7] to visualize hierarchical structures. Although these partitions are not polyhedral subdivisions, they can be constructed from a recursively regular subdivision applying a weighting scheme as in the proof of Theorem 3.7 [24].

The problem of embedding directional graphs is left in a similar situation. A natural and easy to state open question is whether deciding if a directional graph can be embedded in a given point set is NP-hard.

Concerning algorithmic issues, we have proven that the finest regular coarsening and the regularity tree of a subdivision can be computed in polynomial time. We have used these facts to prove that recursive regularity of a subdivision can be decided in polynomial time as well, which is relevant for the algorithmic version of the aforementioned problems.

References

[1] Oswin Aichholzer, Franz Aurenhammer, Siu-Wing Cheng, N. Katoh, Günter Rote, Michael Taschwer, and Yin-Feng Xu. Triangulations intersect nicely. *Discrete & Computational Geometry*, 16(4):339–359, 1996.

[2] David Alberts. A Generalization of the Hungarian Method for the c-Assignment Problem. Diplomarbeit, Freie Universität Berlin, 1992.

[3] Andrei Asinowski, Tillmann Miltzow, and Günter Rote. Quasi-parallel segments and characterization of unique bichromatic matchings. *Journal of Computational Geometry*, 6:185–219, 2015.

[4] F. Aurenhammer, F. Hoffmann, and B. Aronov. Minkowski-type theorems and least-squares clustering. *Algorithmica*, 20:61–76, 1998.

[5] Franz Aurenhammer. A criterion for the affine equivalence of cell complexes in $\mathbb{R}^d$ and convex polyhedra in $\mathbb{R}^{d+1}$. *Discrete & Computational Geometry*, 2(1):49–64, 1987.

[6] Franz Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing*, 16(1):78–96, February 1987.

[7] M. Balzer and O. Deussen. Voronoi treemaps. In *IEEE Symposium on Information Visualization, INFOVIS'05*, pages 49–56. IEEE Computer Society, Oct 2005.

[8] Louis J. Billera and Bernd Sturmfels. Fiber polytopes. *Annals of Mathematics*, 135(3):527–549, 1992.

[9] Prosenjit Bose, Leonidas J. Guibas, Anna Lubiw, Mark H. Overmars, Diane L. Souvaine, and Jorge Urrutia. The floodlight problem. *International Journal of Computational Geometry and Applications*, 7(1/2):153–163, 1997.

[10] Matthew Cary, Atri Rudra, Ashish Sabharwal, and Erik Vee. Floodlight illumination of infinite wedges. *Computational Geometry: Theory and Applications*, 43(1):23–34, January 2010.

[11] Robert Connelly. Rigidity and energy. *Inventiones Mathematicae*, pages 11–33, 1982.

[12] Henry Crapo and Walter Whiteley. Spaces of stresses, projections and parallel drawings for spherical polyhedra. *Contributions to Algebra and Geometry*, 35(2):259–281, 1994.

[13] Jurek Czyzowicz, Eduardo Rivera-Campo, and Jorge Urrutia. Optimal floodlight illumination of stages. In *Canadian Conference on Computational Geometry*, pages 393–398. University of Waterloo, 1993.

[14] Nicolaas G. de Bruijn. Problems 17 and 18. *Nieuw Archief voor Wiskunde*, 2:67, 1954. Answers in *Wiskundige Opgaven met de Oplossingen* 20:19–20, 1955.

[15] Jesus A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations*. Springer, 2010.

[16] Jana Dietel, Hans-Dietrich Hecker, and Andreas Spillner. A note on optimal floodlight illumination of stages. *Information Processing Letters*, 105(4):121–123, 2008.

[17] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Berlin, Heidelberg, 1987.

[18] Herbert Edelsbrunner. An acyclicity theorem for cell complexes in d dimension. *Combinatorica*, 10(3):251–260, 1990.

[19] V. Galperin and G. Galperin. Osveshchenie ploskosti prozhektorami (Illumination of the plane with floodlights, in Russian). *Kvant*, 11:28–30, 1981.

[20] J. Galtier, F. Hurtado, M. Noy, S. Perennes, and J. Urrutia. Simultaneous edge flipping in triangulations. *International Journal of Computational Geometry and Applications*, 13(2), April 2003.

[21] I.M. Gelfand, M. Kapranov, and A. Zelevinsky. *Discriminants, Resultants, and Multidimensional Determinants*. Modern Birkhäuser Classics. Birkhäuser Boston, 2008.

[22] Leo J. Guibas and Frances F. Yao. On translating a set of rectangles. In F. P. Preparata, editor, *Computational Geometry*, volume 1 of *Advances in Computing Research*, pages 61–77. JAI Press, Greenwich, Conn., 1983.

[23] Hiro Ito, Hideyuki Uehara, and Mitsuo Yokoyama. NP-Completeness of stage illumination problems. In Jin Akiyama, Mikio Kano, and Masatsugu Urabe, editors, *Discrete and Computational Geometry*, volume 1763 of *Lecture Notes in Computer Science*, pages 158–165. Springer-Verlag, 2000.

[24] Rafel Jaume. *Tessellations for Geometric Optimization*. PhD thesis, Arbeitsgruppe Theoretische Informatik, Freie Universität Berlin, 2014.

[25] Rafel Jaume and Günter Rote. The finest regular coarsening and recursively-regular subdivisions. arXiv:1310.4372v2 [cs.CG], 2014.

[26] Lionel Pournin. Lifting simplicial complexes to the boundary of convex polytopes. *Discrete Mathematics*, 312(19):2849–2862, 2012.

[27] Günter Rote. Two applications of point matching. In *Proc. 25th European Workshop on Computational Geometry*, Brussels, Belgium, March 2009.

[28] B. Roth and W. Whiteley. Tensegrity frameworks. *Transactions of the American Mathematical Society*, 265(2):419–446, 1981.

[29] Francisco Santos. Non-connected toric Hilbert schemes. *Mathematische Annalen*, 332:645–665, 2005.

[30] Fumihiko Takeuchi. Nonregular triangulations, view graphs of triangulations, and linear programming duality. In Jin Akiyama, Mikio Kano, and Masatsugu Urabe, editors, *Discrete and Computational Geometry*, volume 2098 of *Lecture Notes in Computer Science*, pages 330–338. Springer-Verlag, 2001.

[31] Jorge Urrutia. Art Gallery and Illumination Problems. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, chapter 22, pages 973–1027. North-Holland, Amsterdam, 2000.

A An acyclic polyhedral fan which is not universal

We show here that the polyhedral fan $\mathcal{F}$ described in Proposition 3.16 and shown in Figure 5, which is not universally covering, is acyclic.

The fan $\mathcal{F}$ is a pointed cone with apex O over a two-dimensional subdivision $\mathcal{S}$ in the plane $z = -1/8$. We will show that it is therefore sufficient to look for cycles of the in-front relation in the planar subdivision $\mathcal{S}$.

We choose any point $v \neq 0$ in space (possibly at infinity) and look at the order in which rays ρ emanating from v intersect the cones of $\mathcal{F}$. We are only interested in the part ρ^- of such a ray in the halfspace $z < 0$, and we perform a central projection from O onto the plane $z = -1/8$. The projection ρ' of ρ^- may by a line segment or a ray; in any case, it is part of a ray starting at the projected point v'. (If v lies in the plane $z = 0$, then v' is at infinity, and the projections ρ' are parallel rays.) When we compare the intersection order of ρ with the order in which the projected ray ρ' intersects the cells of $\mathcal{S}$, we either get the

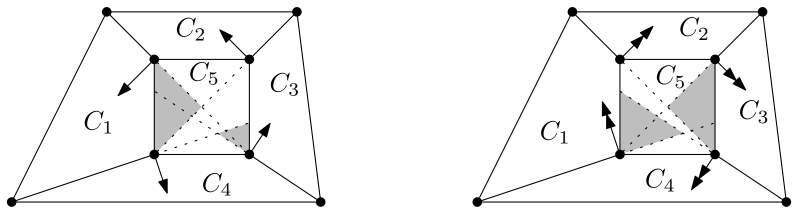

Figure 8: Halfplanes with empty intersection.

same order, if v lies in the lower halfspace $z \leq 0$, or the reverse order, if v lies in the upper halfspace $z > 0$. In either case, we can get a cycle from v only if we get a cycle from v' in the plane subdivision $\mathcal{S}$.

Let us check that $\mathcal{S}$ is acyclic from any point v'. $\mathcal{S}$ is a convex subdivision in the plane. It is known that a convex subdivision in the plane cannot be cyclic from a point at infinity [14, 22], see also [3, Section 3.2, p. 201]. By the same argument that is used to prove this statement, a cycle of the in-front relation around a finite point v' must be formed from cells that *surround* v', in the sense that every ray from v' to infinity must hit one of these cells. Suppose that the point v' lies in some cell C. In this case, all rays emanating from v' meet the interior of C right from the start if they meet C at all. It follows that C is a minimal element for the in-front relation and cannot be part of any cycle.

Now, if v' lies in one of the cells $C = C_1, C_2, C_3, C_4$ that are incident to the outer boundary, there is always a ray from v' that intersects only C. It follows from the above considerations that there can be no cycle around v'. The same argument applies when v' lies in the unbounded region outside $|\mathcal{S}|$. Thus, the only remaining possibility of a cycle is that v' lies in C_5, and the cycle is $C_1C_2C_3C_4$ or its reverse. These two cycles can be easily excluded by looking at the arrangement of the lines separating successive cells in the cycle: If $C_1C_2C_3C_4$ should be a cycle in the in-front relation, v' would have to lie in the four halfplanes indicated by arrows in Figure 8 (left). These halfplanes have empty intersection since the two gray regions in the figure are disjoint. If $C_4C_3C_2C_1$ should be a cycle, v' would have to lie in the four halfplanes indicated by double arrows in Figure 8 (right). In either case, a direct inspection reveals that the four halfplanes have an empty intersection. This finishes the proof of acyclicity.

B An acyclic triangulation which is not recursively regular

We prove here that the triangulation $\mathcal{S}$ pictured in Figure 2 is not recursively regular. To this end, we prove that its finest regular coarsening has only one cell, by showing that its regularity system has no solution.

The rows of the matrix of the regularity system of $\mathcal{S}$ associated to the edges labeled

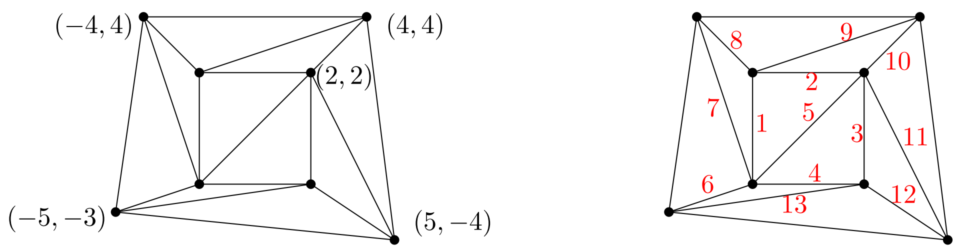

Figure 9: A triangulation which is not recursively regular. The four inner vertices form an axis-aligned 4×4 square.

in Figure 9 are

$$\begin{aligned}
s_1 &= (8, -32, 8, 0, 16, 0, 0, 0)\\
s_2 &= (8, 0, -24, 0, 0, 16, 0, 0)\\
s_3 &= (12, 0, 8, -36, 0, 0, 16, 0)\\
s_4 &= (-28, 0, 4, 8, 0, 0, 0, 16)\\
s_5 &= (-16, 16, -16, 16, 0, 0, 0, 0)\\
s_6 &= (-48, 0, 0, 20, 4, 0, 0, 24)\\
s_7 &= (-16, 20, 0, 0, -12, 0, 0, 8)\\
s_8 &= (16, -48, 0, 0, 24, 8, 0, 0)\\
s_9 &= (0, -16, 16, 0, 8, -8, 0, 0)\\
s_{10} &= (0, 18, -50, 0, 0, 24, 8, 0)\\
s_{11} &= (0, 0, -22, 18, 0, 12, -8, 0)\\
s_{12} &= (0, 0, 17, -57, 0, 0, 28, 12)\\
s_{13} &= (17, 0, 0, -13, 0, 0, 4, -8).
\end{aligned}$$

We form a linear combination of the corresponding inequalities with the following positive coefficients y_i:

$$y_1 = 207, y_2 = 24, y_3 = 20, y_4 = 24, y_5 = 288, y_6 = 24, y_7 = 1308,$$
$$y_8 = 24, y_9 = 1464, y_{10} = 24, y_{11} = 198, y_{12} = 24, y_{13} = 1464.$$

Since $\sum_{i=1}^{13} y_i s_i = 0$, this linear combination leads to the contradiction $0 > 0$, and hence the regularity system is inconsistent.

We prove now that $\mathcal{S}$ has no cycle in the in-front relation. We show first that the coarsening $\mathcal{S}'$ of $\mathcal{S}$ depicted in Figure 10 is acyclic. This is can be seen, as in Appendix A, by checking that the shaded regions in the figure are disjoint. On the other hand, the subdivisions $\mathcal{S}|_C$ for $C \in \mathit{cells}(\mathcal{S}')$ are obviously acyclic. Hence, a cycle cannot be contained in $\mathcal{S}|_C$ for $C \in \mathit{cells}(\mathcal{S}')$, but it also cannot involve cells of $\mathcal{S}'$ contained in different cells of $\mathcal{S}'$. Thus, the subdivision $\mathcal{S}$ is acyclic.

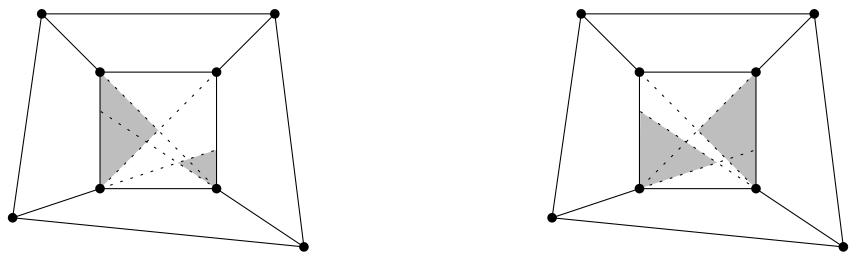

Figure 10: An acyclic coarsening $\mathcal{S}'$ of $\mathcal{S}$.

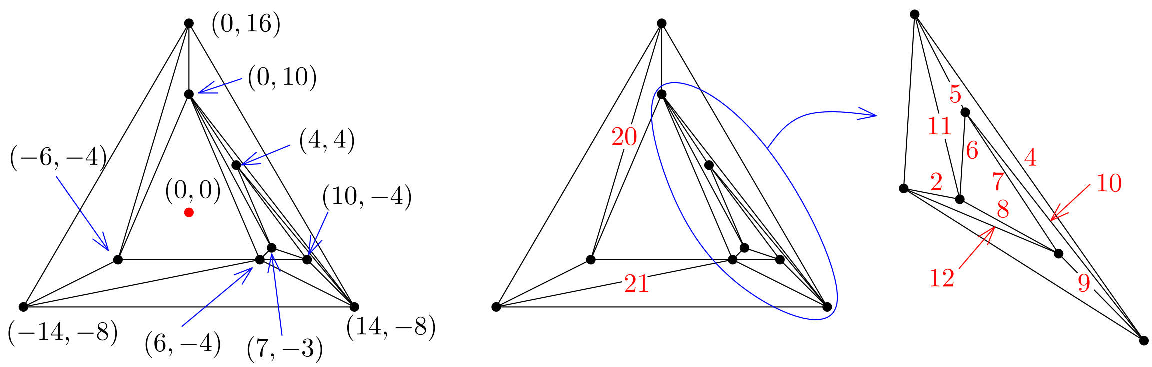

Figure 11: A recursively regular subdivision which is not regular.

C A subdivision whose regularity tree has height two

We prove here that the subdivision $\mathcal{S}$ pictured in Figure 1 (left) is not regular and that its finest regular coarsening is the subdivision $\mathcal{S}_0$ defined by the second level of the tree in Figure 1 (right).

The rows of the matrix of the regularity system of $\mathcal{S}$ associated to the edges labeled with numbers in Figure 11 are

$$
\begin{aligned}
s_2 &= (-56, 20, 0, 4, 32, 0, 0, 0, 0)\\
s_4 &= (0, 0, 84, -72, 0, -24, 0, 0, 12)\\
s_5 &= (12, 0, -56, 34, 0, 10, 0, 0, 0)\\
s_6 &= (4, 10, -32, 18, 0, 0, 0, 0, 0)\\
s_7 &= (8, -34, 8, 0, 0, 18, 0, 0, 0)\\
s_8 &= (-32, 10, 4, 0, 18, 0, 0, 0, 0)\\
s_9 &= (0, 56, 16, 0, 8, 32, 0, 0, 0)\\
s_{10} &= (0, 12, -16, 8, 0, -4, 0, 0, 0)\\
s_{11} &= (-20, 0, 20, -10, 10, 0, 0, 0, 0)\\
s_{12} &= (16, -12, 0, -8, 4, 0, 0, 0, 0)
\end{aligned}
$$

$$s_{20} = (0, 0, 0, 136, 0, 0, 36, -84, -88)$$
$$s_{21} = (0, 0, 0, 0, -112, 48, -48, 112, 0).$$

We form a linear combination of the inequalities with the following nonnegative coefficients y_i:

$$y_2 = \tfrac{1}{10}, y_4 = \tfrac{11}{10}, y_5 = \tfrac{109}{110}, y_6 = 1, y_7 = \tfrac{50}{99}, y_8 = \tfrac{71}{99},$$
$$y_9 = \tfrac{1}{10}, y_{10} = \tfrac{11}{10}, y_{12} = \tfrac{4}{5}, y_{11} = \tfrac{23}{110}, y_{20} = \tfrac{3}{20}, y_{21} = \tfrac{9}{80}.$$

The remaining values y_i are zero. Since $\sum_{i=1}^{21} y_i s_i = 0$, this linear combination leads to the contradiction $0 > 0$, and hence the regularity system is inconsistent.

The sub-subdivision of $\mathcal{S}$ represented in the lower part of Figure 11 is a variant of a typical non-regular subdivision appearing, for instance, in [15]. However, since the lines supporting the edges 2, 5 and 9 are concurrent, the sub-subdivision is recursively regular (its finest regular coarsening is the projection of a truncated triangular pyramid). Similarly, it is clear that $\mathcal{S}_0$ is regular as well.