

An Efficient Local Approach to Convexity Testing of Piecewise-Linear Hypersurfaces

Konstantin Rybnikov

*One University Ave., Olney Hall 428, University of Massachusetts at Lowell,
Lowell, MA 01854 USA*

Email address: Konstantin.Rybnikov@uml.edu

Abstract

We prove the following *criterion*: a compact connected piecewise-linear hypersurface (without boundary) in \mathbb{R}^n ($n \geq 3$) is the boundary of a convex body if and only if every point in the relative interior of each $(n-3)$ -face has a neighborhood that lies on the boundary of some convex body. This criterion is derived from our theorem that any connected complete locally-convex hypersurface in \mathbb{S}^n ($n \geq 3$) is the boundary of a convex body in \mathbb{S}^n . We give an easy-to-implement convexity testing algorithm based on our criterion. This algorithm does not require any assumptions about the global topology of the input hypersurface. For \mathbb{R}^3 the number of arithmetic operations used by our algorithm is at most linear in the number of vertices, while in general it is at most linear in the number of incidences between $(n-2)$ -faces and $(n-3)$ -faces. The algorithm still remains polynomial even when the dimension n is a variable and the bit complexity model for (exact) arithmetic operations is used. The suggested method works in more general situations than the convexity verification algorithms developed by Mehlhorn et al. (1996) and Devillers et al. (1998) – for example, our method does not require the input surface to be homeomorphic to the sphere, nor does it require the input data to include normal vectors to the facets that are oriented “in a coherent way”. For \mathbb{R}^3 the complexity of our algorithm is the same as that of previously known algorithms; for higher dimensions there seems to be no clear winner, but our approach is the only one that easily treats surfaces of arbitrary topology without a preliminary topological computation aimed at verifying that the input is a topological sphere. Furthermore, our method can be easily extended to piecewise-polynomial surfaces of degree 2 and 3.

Key words: Program checking, Output verification, Geometric property testing, Convexity, Piecewise-linear surface, Polyhedral surface

Saturday, January 5, 2008 07:31

5 January 2008

1 Introduction

Blum and Kannan (1989) suggested a paradigm of output verification. Since a complete check of a program is often difficult or impossible – for example, when the source code has not been made public – it is important to have algorithms that verify key properties of mathematical objects generated by programs. Instead of the source code verification one can try to verify the properties of the output that are deemed essential by users of the program. In computational geometry this paradigm was developed, among others, by Mehlhorn et al. (1996, 1999) and Devillers et al. (1998). For example, the LEDA C++ library contains programs verifying the convexity of a polygon, Delaunay property of a tiling, etc (Mehlhorn and Näher, 2000). Devillers et al. argue that it is easier to evaluate the quality of the output of a geometric algorithm, than the correctness of the algorithm or program producing it. This paper contributes to the problem of verification of convexity of a large class of piecewise-linear (PL) hypersurfaces in \mathbb{R}^n for $n \geq 3$. The novelty of our approach is in reducing the verification of global convexity of a PL-hypersurface to the verification of local convexity at the faces of small codimension.

We show that a compact connected PL-hypersurface (without boundary) realized in \mathbb{R}^n ($n \geq 3$) without local self-intersections is the boundary of a convex body if and only if the relative interior of each $(n - 3)$ -face has a point such that a small Euclidean ball centered at this point is cut by the hypersurface into two pieces, one of which is convex. This *local convexity* condition can also be expressed as that the point has a neighborhood (on the hypersurface), which lies on the boundary of a convex body; such a point is called a point of local convexity.

Dropping the compactness requirement invalidates the above criterion. For example, in \mathbb{R}^3 the direct (affine) product of a *non-convex* simple 4-gone in the xy -plane and a line not collinear to the xy -plane, does not bound any convex body (Figure 1 shows a part of such an unbounded surface). Fortunately, the criterion can be “repaired” to include unbounded




Fig. 1. The product of a non-convex 4-gone and a line is locally convex, but not globally convex hypersurfaces. If we require that the hypersurface, in addition to the conditions mentioned in the above paragraph, has at least one point of *strict convexity*, then the local convexity at the

$(n-3)$ -faces implies global convexity. A point s on the hypersurface is called a *point of strict convexity* if it is a point of local convexity and, additionally, there is a small ball centered at s such that its intersection with the hypersurface, for the exception of point s , lies in an *open halfspace* with respect to some hyperplane through s . For a PL-hypersurface a point of strict convexity is always a vertex. While the technical terms are defined in Section 2, we would like to point out that the class of surfaces covered by our convexity theorem includes all simplicial hypersurfaces (without boundary) whose intersection with any bounded subset of \mathbb{R}^n involves only a finite number of simplices. One who is primarily interested in simplicial surfaces in \mathbb{R}^3 (also known as “triangular meshes” in applications) can skip Section 2 at the first reading. To get an intuitive grasp of the situation one can simply think of PL-surfaces (without boundary) in \mathbb{R}^3 that are closed as subsets of \mathbb{R}^3 .

Our treatment of the subject is based on the idea of separation of topological, combinatorial, and geometric properties of a surface under testing. For us, in general, a surface is a triple (**Topology**, **Combinatorics**, **Realization**), where the first attribute describes an abstract topological space (usually a manifold), the second a partition of this space into cells (e.g. simplices), and the third a map from the topological space into \mathbb{R}^n (or \mathbb{S}^n). Of course, **Topology** leaves its imprint on **Combinatorics**, and **Realization** is usually required to be well-behaved with respect to both **Topology** (at least continuous) and **Combinatorics** (we will insist that each k -dimensional cell is mapped homeomorphically onto a polyhedral subset of affine dimension k). Similar separation of properties is followed in the treatment of local convexity: a hypersurface $r : \mathcal{M} \rightarrow \mathbb{R}^n$, where \mathcal{M} is a manifold, is said to be locally convex at $p \in \mathcal{M}$ if (1) r , restricted to some neighborhood \mathcal{N} of p , is a homeomorphism and (2) $r(\mathcal{N})$ lies on the boundary of a convex body in \mathbb{R}^n . For example, the PL-realization of the octahedral decomposition of the sphere shown in Figure 2 (right) does not satisfy the first (homeomorphism) condition, while obviously satisfying the second. This PL-realization does not bound any convex body. However, if in the same example $r(e)$ happened to coincide with $r(f)$, the image $r(\mathbb{S}^2)$ would be the boundary of a convex pyramid, although in this case r would not be a homeomorphism onto the boundary of this pyramid. These examples show the importance of being rather pedantic in regard to distinguishing between the objects living on the abstract manifold and their images in \mathbb{R}^n . In fact, the right way to think of local convexity is in terms of an abstract manifold equipped with a convexity structure locally, so that the local convexity structures over any two neighborhoods agree on their intersection, just as we think of differential manifolds or algebraic varieties – a reader familiar with the modern view on these theories (which is not at all required for understanding this paper) may now recognize the sheaf-theoretic nature of local convexity. On the other hand, it is quite safe to visualize $r : \mathcal{M} \rightarrow \mathbb{R}^3$ locally as a fragment of a convex Euclidean surface.

Surfaces under consideration are allowed to intersect themselves, but only in such a way that no local singularities appear. The mathematical notion capturing the concept of such a realization is called *immersion*. For example, in Figure 3 the curve on the left has self-intersections; however, this curve can be thought of as an immersion of a regular hexagon (the polygonal curve in the center with the standard topology, i.e. induced by the topology of \mathbb{R}^2) into the plane. On the other hand, the curve on the right has a *local singularity* at the point $r : a \rightarrow r(a)$: no matter how small a neighborhood around a we consider, the map

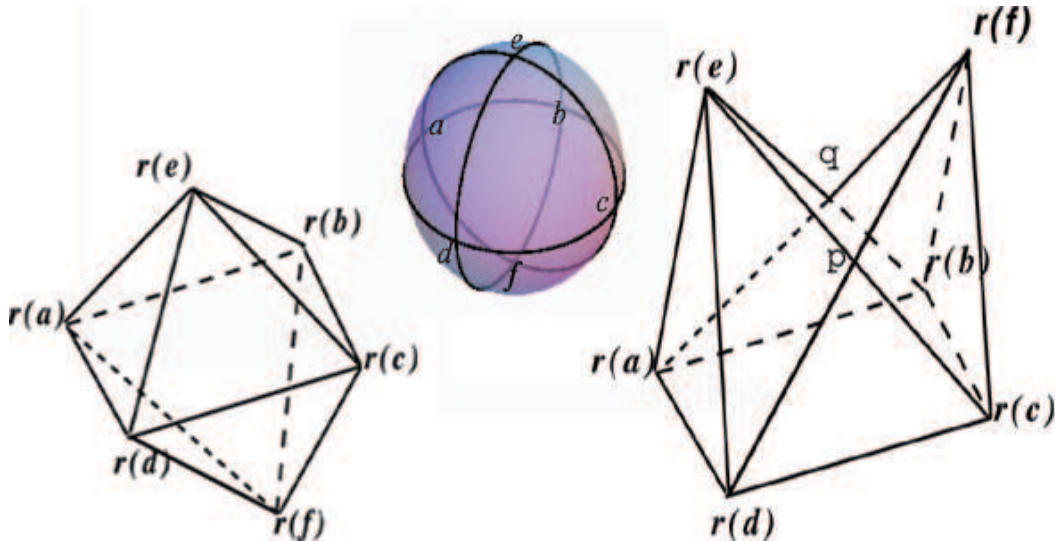


Fig. 2. Center: $(\mathbb{S}^2, \mathcal{O})$, where \mathcal{O} stands for the octahedral partition of \mathbb{S}^2 . Left: convex PL-embedding of $(\mathbb{S}^2, \mathcal{O})$ into \mathbb{R}^3 . Right: realization with self-intersections; here each point of \mathbb{S}^2 is mapped onto the surface of a convex body, however, the local homeomorphism condition is not respected on the equator (a, b, c, d) .

r will not be bijective on it.

It is important to understand that it would not be correct to think of immersion as of a realization such that for each point of the surface there is a ball centered at this point within which the surface does not intersect itself. For example, if we modify 3 (left) by moving $r(\mathbf{f})$ and $r(\mathbf{c})$ towards each other so that eventually $r(\mathbf{f}) = r(\mathbf{c})$ and the curve looks like a bow-tie, the resulting 1-surface is still an immersion.

An immersion which is a global homeomorphism onto the image (which guarantees the absence of any self-intersections) is called an *embedding*. Indeed, a surface that intersects




Fig. 3. The left figure shows an immersion of an abstract hexagon into \mathbb{R}^2 , while the right figure shows an embedding.

itself cannot serve as the boundary of a convex body; however, it is one of the key points of this paper that we can ignore testing for global self-intersections – they are ruled out automatically when the local convexity at the $(n - 3)$ -faces is verified (including the local homeomorphism property).

We will need some notions of combinatorial topology, such as regular (or semiregular – see Section 2) cell-complexes (a.k.a. CW-complexes). This need arises from the natural desire

to consider surfaces which are not necessarily simplicial. This desire stems not only from mathematical curiosity, but also from the demands of potential applications: for example, the CGAL C++ library has a class template *Polyhedron_3*, which represents the combinatorics of a not necessarily simplicial cell-partition of an orientable 2-dimensional PL-manifold with or without boundary (called “polyhedral surface” in CGAL). While for \mathbb{R}^3 the notion of a general PL-surface can be stated in rather simple terms, essentially by replacing (solid) triangles with (solid) simple polygons (see Kettner (1999) for details), for $n > 3$ certain extra care is needed in order to avoid some nasty topological pitfalls. Although the surfaces we consider may seem rather general, in the case of finitely many cells and absence of self-intersections, they form a very special subclass of NEF-polyhedra (see Hachenberger et al (2007) for NEF-polyhedra). Furthermore, using NEF-polyhedra seems to be a natural way for implementing our approach in higher dimensions. Currently CGAL supports 2D and 3D NEF polyhedra, which makes it possible to implement the class of PL-hypersurfaces considered in this paper for $n = 3$ and $n = 4$ (by implementing the geometric realization $r(C)$ of each 3-cell C as a NEF-polyhedron in the affine span of $r(C)$).

In this paper we construct an algorithm for convexity testing that can be applied to any compact PL-hypersurface without boundary. Our approach does not require any knowledge of the global topology of the surface. Even in the spherical case ($\mathcal{M} \cong \mathbb{S}^{n-1}$) the direct comparison of the complexity of our algorithm and those of Mehlhorn et al. (1996b; 1999) and Devillers et al. (1998) is not quite meaningful, since these authors made the following simplifying assumptions.

- (CO) *The surface under testing is an oriented compact hypersurface. The normals to the $(n-1)$ -faces are given as part of the input, and they are all oriented either outwards or inwards (“Coherent Orientation”); in other words, the orientation of the input surface is given geometrically.*
- (S) *The cell-partition is Simplicial.*

It can be shown that (S) is not necessary for correctness of these algorithms, if it is known that the $(n-1)$ -faces are convex. However, the assumption (S) seriously affects the complexity analysis, which the authors performed only for the simplicial case. It is not clear to us why the assumption (CO) is natural; it seems that for $n > 3$ a convex hull “builder” is at a potential risk of producing not only a geometrically non-convex output, but also a non-spherical (combinatorially) complex. For $n \geq 4$ it is not easy to design efficient checking procedures that guarantee that the simplicial (or cell) complex under construction remains spherical. Our approach does not require (CO) and (S). For \mathbb{R}^3 , when (CO) and (S) hold, both our and previous approaches have the same complexity, which is $O(f_0)$, where f_0 stands for the number of vertices in the triangulation.

1.1 Outline of the paper

The remainder of Section 1 highlights the paradigm of local approach in the context of geometric property testing and compares our local approach to previous algorithmic approaches to verification of convexity. Section 2 gives definitions and notation; it also gives the complete

formal description of the main results of our work. The language of the paper presupposes only basic familiarity with partially ordered sets (*posets*), linear algebra and geometry, point set topology, Euclidean convexity theory (e.g. Rockafellar, 1997, Part IV), and combinatorial topology (e.g. Seifert and Threlfall, 1980). The exposition is practically self-contained, for the exception of some basic notions of topology such as the fundamental group and covering mappings. We expect that anybody with basic knowledge of linear algebra and classical Euclidean geometry (e.g. Ch. 1,2, and 4 of Kostrikin and Manin (1989) or equivalent) will be able to understand the paper except, maybe, for some topological arguments in proofs given in Section 4. We have made every effort to ensure that the pseudocode is accurate and complete and we hope that it can be used *as is* by any reader familiar with standard linear algebra and geometry. Section 3 explains the local approach to convexity verification for general (not necessarily PL) hypersurfaces. At the end of that section we explain why the study of convexity properties of PL-surfaces in \mathbb{R}^n necessitates the study of convexity properties of surfaces in the spherical space \mathbb{S}^n . There we also sketch the proof of our main theorem. Section 4 is devoted to convexity testing in \mathbb{S}^n . In Section 5 we prove our main Theorem 13. Section 6 gives the algorithm and proves its correctness; Section 7 is devoted to the complexity analysis. Finally, Section 8 summarizes the paper.

1.2 Previous results on convexity verification and advantages of the new approach

We emphasize a *purely local approach* to the verification of convexity. Here are the reasons for advocating the local approach.

Local convexity implies global convexity: As shown by Van Heijenoort (1952) and Jonker & Norman (1973), under rather mild general assumptions, the global convexity verification for a closed immersed hypersurface reduces to local convexity verification. Thus, at least philosophically, it makes sense to have a local algorithmic approach to convexity testing. Simply put, this is the right way to think. The earlier convexity verification algorithms made use of the concepts of *core* (Mehlhorn et al.) and *seam* (Devillers et al.) of a polyhedral hypersurface, which are both defined globally and *only* for a hypersurface with a given “coherent” field of outer (or inner) facet normals. It should be noted that the previous authors also used the words **local convexity** and applied it to $(n - 2)$ -dimensional faces. However, their usage of **local** was somewhat misleading, since their definition made use of an already existing global orientation of the hypersurface, which was presumed *a priori* given. From the truly local point of view any PL-hypersurface is locally-convex at its $(n - 2)$ -faces. The concept of local convexity used in this paper was introduced by Van Heijenoort and is free of any global assumptions. When our algorithm finds *a* violation of local convexity it reports the type of violation, i.e. whether it is a violation of the immersion property (in the PL-case this always boils down to violation of the local injectivity) or a violation of convexity, and the $(n - 3)$ -face whose star failed to be locally convex. This information is the *certificate of violation*.

Convexity checking of surfaces changing in time: An important practical consideration is repeated convexity verification of a surface that is being gradually modified over time. For simplicity, consider the case of a simplicial convex surface in \mathbb{R}^3 . If it is known

that each modification affects only a small number of vertices, then, according to our main Theorem 13, we only have to recheck for convexity at the affected vertices. In the language of kinetic data structures (Guibas, 2004, p. 1121) our convexity certification is *local*. On the other hand, it might not be so easy to ramify the algorithms of Mehlhorn et al. and Devillers et al. so that the time required for rechecking for convexity after a modification would be commensurate with the (combinatorial) size of the modification. Note that the geometric magnitude of modifications does not affect the complexity of rechecking in the real RAM model of computation. Even under the exact computing paradigm, if the geometric magnitude of modifications is within reasonable limits (e.g. the perturbed coordinates are within a constant factor of the original ones), the complexity of rechecking is bounded by a constant times the combinatorial size of the modification. In other words, our algorithm shows that the complexity of convexity verification depends “continuously” on the combinatorial size of the modification.

Topological generality: Our algorithm works in the same way for both spherical and non-spherical surfaces. To the contrary, the previous papers on the subject assume that the input surface is spherical and this assumption is used in the correctness proofs presented in Devillers et al. (1998) and Mehlhorn et al. (1999) (we do not know if these algorithms actually remain correct for non-spherical inputs). Note that for even n one cannot distinguish between the sphere \mathbb{S}^{n-1} and other compact $(n - 1)$ -manifolds on the basis of Euler characteristic alone, as all such manifolds have Euler characteristic of zero. Note that for $n \geq 4$ the problem whether a given $(n - 1)$ -manifold, defined combinatorially, is \mathbb{S}^{n-1} is highly non-trivial (recall Poincaré’s conjecture, the problem of classification of 3-manifolds, etc). Note that classical homology groups and Euler characteristic are not able to discern between a sphere and a homology sphere (see Seifert, Threlfall, 1980). We also notice that our theorem and algorithm hold for homology manifolds.

Convex surfaces with non-convex faces. Our approach is the only one that treats surfaces with non-convex facets without the messy (and difficult for $n \geq 4$) task of triangulating them. Note that a PL-hypersurface may have non-convex facets and still be convex (see Figure 4).




Fig. 4. PL-surface with non-convex faces. Geometrically it is the boundary of a tetrahedron.

No need for large-scale linear algebraic and linear programming computations: The amount of linear-algebraic computations for *non-simplicial* inputs is significantly smaller in our algorithm than in the previous algorithms. In particular, our approach completely avoids linear programming, in all dimensions. For example, the algorithm of

Mehlhorn et al. (1999; Sec. 2.1), if adopted for surfaces with convex, but not necessarily simplicial facets, requires checking, for each facet of the hypersurface, whether a certain ray belongs to the cone built over this facet (facet cone). That would require either linear programming or obtaining a complete description of all facet cones in terms of inequalities.

Extensibility: Our approach extends to spline (piecewise-polynomial) surfaces of small degree (2 or 3) in \mathbb{R}^3 and \mathbb{R}^4 (this work is in progress). For example, when a surface in \mathbb{R}^3 is made of polynomial patches, where each patch is defined as an explicit quadratic map from a triangle to \mathbb{R}^3 , testing for convexity at each vertex is not difficult.

Efficiency: A detailed complexity analysis is given at the end of the paper. For \mathbb{R}^3 all algorithms have the same complexity and should have comparable performance when used on *simplicial surfaces* of spherical topology. Previous papers on the subject do not really discuss non-simplicial surfaces, although we have verified that all of their convexity criteria hold for PL-hypersurfaces of *spherical topology* with *convex facets*. Our algorithm is more efficient than that of Devillers et al. Let us denote by $f_{i,j}$ the number of incidences between the i -faces and the j -faces; we abbreviate $f_{i,i}$ as f_i . While Devillers et al. (1998) stated $O(f_0)$ as the running time for any dimension, this was clearly a typo: this bound is impossible, even for PL-spheres, by fundamental counting theorems of polyhedral combinatorics; furthermore, the pseudocode in Devillers et al. (1998: Section 3) has running time for surface Γ , in the notation of that paper, of $\sum_{j=0}^{d-2} f_j + \sum_{\{F \in \Gamma \mid \dim F = d-2\}} \sum_{j=3}^d f_{j-1} j(F)$. For fixed dimension and simplicial inputs our algorithm has the same complexity of $O(f_{n-1})$ as that of Mehlhorn et al. For variable dimension (as in the complexity theory of linear programming) and simplicial inputs our algorithm has the same worst case complexity of $O(n^3 f_{n-1})$ as that of Mehlhorn et al. Note that in our approach the stars of all corners and ridges can be processed in parallel. With the number of processors that scales linearly with $f_{n-2} f_{n-3}$, our approach has $O(1)$ complexity.

Numerical Robustness: Suppose the dimension n is a variable, or just large. Let the input surface be (combinatorially) simplicial and given by the poset of corners, ridges, and facets together with the vertex coordinates. If the input is *geometrically* generic enough so that for each corner no three ridges from the star of this corner lie on the same hyperplane, then our algorithm requires no divisions, but only evaluation of polynomial predicates of degree at most 3.

To summarize, our algorithm is the algorithm of choice for non-simplicial hypersurfaces and in the cases where the topology of the underlying manifold cannot be easily verified. Mehlhorn et al.'s algorithm seems to be the easiest to implement on the basis of standard linear-algebraic routines for simplicial spherical surfaces in \mathbb{R}^3 , provided the outer (or inner) facet normals are provided. Of all the algorithms that we have discussed Mehlhorn et al.'s one is the most global in nature, while Devillers et al.'s algorithm occupies an intermediate position. Devillers et al.'s algorithm does a number of local checks on global surfaces (i.e. determined by the entire input hypersurface) of smaller dimensions iteratively derived from the input hypersurface; these derived surfaces are global in the sense that they are determined by the entire input surface. In contrast, our algorithm checks a single cone in \mathbb{R}^3 derived from the star of each $(n-3)$ -face. It would be interesting to combine the algorithm of Devillers

et al. with our algorithm. Such a hybrid is not likely to have a better worst case complexity, but may have a superior practical performance for surfaces in \mathbb{R}^3 .

2 Definitions, notation, and results

From now on \mathbb{X}^n (or just \mathbb{X}) denotes \mathbb{R}^n or \mathbb{S}^n . As an input surface may have self-intersections, it is convenient to distinguish between an abstract $(n - 1)$ -manifold and its realization in \mathbb{X}^n .

Definition 1 *A hypersurface in \mathbb{X}^n is a continuous realization map $r : \mathcal{M} \rightarrow \mathbb{X}^n$, where \mathcal{M} is a manifold of dimension $n - 1$, with or without boundary.*

Hereafter $r|_{\mathcal{S}}$ denotes the restriction of r to a subset \mathcal{S} of \mathcal{M} , which we also write as $r : \mathcal{S} \rightarrow \mathbb{X}^n$. We will often use notation (\mathcal{M}, r) to refer to a hypersurface. Whenever we want to include manifolds with boundary into our considerations we explicitly say so. To avoid a common confusion caused by (at least three) different usages of **closed** in geometry-in-the-large, we use this word for closed subsets of topological spaces only. We will not use the term “closed surface” at all; a closed submanifold stands for a submanifold which happens to be a closed subset in the ambient manifold.

$i : \mathcal{M} \rightarrow \mathbb{X}$ is called an *immersion* if i is a local homeomorphism; we also refer to (\mathcal{M}, i) as a hypersurface immersed into \mathbb{X} . This is a common definition of immersion in the context of non-smooth geometry-in-the-large; a more restrictive definition is used in differential geometry, furthermore, some authors define immersion as a continuous local bijection. Although the latter definition is not, in general, equivalent to the common one, it is equivalent to it in the context of the theorems stated in this paper.

$e : \mathcal{M} \rightarrow \mathbb{X}$ is called an *embedding* if e is a homeomorphism onto $e(\mathcal{M})$. Obviously, an embedding is an immersion, but not vice versa.

A set $K \subset \mathbb{X}^n$ is called *convex* if for any $x, y \in K$ there is a geodesic segment of minimal length with end-points x and y that lies in K . A *convex body* in \mathbb{X}^n is a closed convex set of full dimension; *a convex body may be unbounded*. The hypersurface (\mathcal{M}, r) is called *locally convex at $p \in \mathcal{M}$* if we can find a neighborhood $\mathcal{N}_p \subset \mathcal{M}$ and a convex body $K_p \subset \mathbb{X}^n$ for p such that $r|_{\mathcal{N}_p} : \mathcal{N}_p \rightarrow r(\mathcal{N}_p)$ is a homeomorphism and $r(\mathcal{N}_p) \subset \partial K_p$. In such a case we refer to K_p as a *convex witness for p* . Thus, the local convexity at $\mathbf{p} = r(p)$ may fail because r is not a local homeomorphism at p or because no neighborhood \mathcal{N}_p is mapped by r onto the boundary of a convex body. In the first case we say that the immersion assumption is violated, while in the second case we say that the convexity is violated. (Here, as everywhere else, the subscript indicates that K_p depends on p in some way but is not necessarily determined by p uniquely.) Often, when it is clear from the context that we are discussing the properties of r near $\mathbf{p} = r(p)$, we say that r is *convex at \mathbf{p}* . If K_p can be chosen so that $K_p \setminus r(p)$ lies in an open half-space defined by some hyperplane passing through $r(p)$, the realization r is called *strictly convex at p* . *From now on we will often apply local techniques of Euclidean convex geometry to \mathbb{S}^n without restating them explicitly for the spherical case.*

Let us recall (see e.g. Rockafellar, 1997) that a point \mathbf{p} on the boundary of a convex set C is called *exposed* if C has a support hyperplane that intersects \bar{C} , the closure of C , only at \mathbf{p} . Thus, an *exposed* point on a convex body K is a *point of strict convexity* on the hypersurface ∂K . Conversely, for a point of strict convexity $p \in \mathcal{M}$ for (\mathcal{M}, r) the image $i(p)$ is an exposed point of any convex witness for p . Local convexity can be defined in many other, non-equivalent, ways (e.g., see van Heijenoort).

A hypersurface (\mathcal{M}, r) is (globally) *convex* if there exists a convex body $K \subset \mathbb{X}^n$ such that r is a homeomorphism onto ∂K . Hence, we exclude the cases where $r(\mathcal{M})$ is the boundary of a convex body, but r fails to be injective. Our algorithm for PL-hypersurfaces will always detect a violation of the immersion property; when $r(\mathcal{M})$ is the boundary of a convex body, but r is not a homeomorphism, the algorithm will produce the negative answer without trying to determine if $r(\mathcal{M})$ is the boundary of a convex body. Of course, the algorithmic and topological aspects of this case may be interesting to certain areas of geometry, such as origami. Note that for $n \geq 3$ an $(n - 1)$ -manifold \mathcal{M} without boundary cannot be immersed into \mathbb{R}^n by a *non-injective* map r so that $r(\mathcal{M})$ is the boundary of a convex set, since any convex hypersurface in \mathbb{R}^n is simply-connected and any immersion onto a simply-connected manifold must be a homeomorphism (this is a consequence of the *covering mapping theorem* in topology – see e.g. Seifert and Threlfall, 1980). However, such immersions cannot be ruled out in the hyperbolic space \mathbb{H}^n , as there are infinitely many topological types of connected convex hypersurfaces in \mathbb{H}^n for $n > 2$ (Kuzminykh, 2005) and some of them are not simply-connected.

By a *subspace* of \mathbb{X}^n we mean an *affine* subspace (i.e. a subspace defined by a system $A\mathbf{x} = \mathbf{b}$) in the case of \mathbb{R}^n , and the intersection of $\mathbb{S}^n \subset \mathbb{R}^{n+1}$ with a *linear* subspace of \mathbb{R}^{n+1} in the case of \mathbb{S}^n . We often use *k-subspace* (or *k-plane*) instead of *k-dimensional subspace*; same convention applies to *k-faces*, *k-cells*, *k-polytopes* etc. A *hyperplane* is a subspace of \mathbb{X}^n of codimension one; a line is a subspace of dimension one. If $\mathcal{S} \subset \mathcal{M}$ is a submanifold with or without boundary, then $\dim \mathcal{S}$ stands for its dimension. For any $S \subset \mathbb{X}^n$ we denote by $\dim S$ the dimension of the minimal subspace containing S ; when this subspace is unique it is denoted by $\text{aff } S$. Thus, in general, $\dim \mathcal{S} \neq \dim r(\mathcal{S})$.

We mostly focus on convexity of PL-hypersurfaces, in particular, boundaries of polytopes. Since a PL-surface under testing may have self-intersections, we cannot just identify the faces of this surface with subsets of \mathbb{R}^n as it is done in the study of convex polyhedra. Arguably, immersions with self-intersections form a majority of cases where convexity fails in “real applications”. We resolve a potential ambiguity by separating the abstract combinatorial information about the cell-partition of \mathcal{M} and the geometric information, such as the positions of the vertices, that describes the realization of \mathcal{M} in \mathbb{R}^n . Since we study PL-realizations, we insist that the realization map r respects the cell-structure on \mathcal{M} , i.e., r maps each *k-cell* $C \subset \mathcal{M}$ to a relatively open *k-polytope* $r(C)$ in \mathbb{R}^n ; the set $r(C)$ need not be convex, but must be topologically equivalent to an open *k-ball*.

Denote by B^d the closed unit ball in \mathbb{R}^d . A denumerable (disjoint) partition \mathcal{P} of a topological space \mathcal{M} is called a *semiregular cell-partition* if

- (1) each $C \in \mathcal{P}$, called a cell of $(\mathcal{M}, \mathcal{P})$, is assigned a number $\dim C \in \mathbb{N}$ ($\dim C \leq \dim \mathcal{M}$), called the dimension of C , so that C is homeomorphic to $\text{int } B^{\dim C}$, the interior of $B^{\dim C}$;
- (2) the closure \overline{C} (in \mathcal{M}) of each $C \in \mathcal{P}$ is the union of C and cells of smaller dimensions;
- (3) for each $C \in \mathcal{P}$ there is a mapping $r_{\overline{C}} : \overline{C} \rightarrow B^d$, where $d = \dim C$, which is a homeomorphism onto $r_{\overline{C}}(\overline{C})$ and whose restriction to C is a homeomorphism onto the interior of B^d .

With any semiregular cell-partition there is a natural structure of poset. Namely, for cells F and C we have $F \preceq C$ if and only if $F \subset \overline{C}$, where \overline{C} stands for the closure of C in the topology of \mathcal{M} . If $F \preceq C$, we say that F is a face of C . We will use the same symbol \mathcal{P} for the partition and its poset and we do not distinguish in print between C as an element of \mathcal{P} and C as a subset of \mathcal{M} . Any semiregular cell-partition is an *abstract polytope* in the sense of McMullen and Schulte. When $(\mathcal{M}, \mathcal{P})$ can be thought of as the surface of a convex polytope P , the poset \mathcal{P} , augmented with an infimum (\emptyset) and a supremum (symbolizing all of P), is known as the *face-lattice* of P . In our, more general, case the augmented \mathcal{P} *may fail to be a lattice*, since in a regular cell-partition two k -cells may share more than one maximal face and serve as maximal faces to more than one cell of higher dimension. We use $\text{Sk}_d(\mathcal{M}, \mathcal{P})$ to denote the d -skeleton of $(\mathcal{M}, \mathcal{P})$ - i.e. the subcomplex of $(\mathcal{M}, \mathcal{P})$ that consists of all cells of dimension d or less. $\text{Star } F$ denotes the union of all cells whose closures contain F (we use $\text{Star } F$ to denote both the subset of \mathcal{M} and the corresponding sub-poset of \mathcal{P}). If each cell of \mathcal{P} is contained in only finitely many closures of cells of \mathcal{P} , the partition \mathcal{P} is called *star-finite*. If the closure of each cell is the union of finitely many cells, the partition is called *closure-finite*. When a partition is both closure- and star-finite, it is called *locally-finite*. From now on we will consider *only star-finite cell-partitions*.

Cell-partitions are also known as topological cell-complexes in literature. Often in the definition of a cell-partition one insists that the closure of each cell is the image of a closed ball, which forces the compactness for the closure of each cell – we do not make such a requirement. Hence, by our definition the closures of cells can be “semiclosed-semiopen”. Our notion of semiregular cell-partition is a natural generalization of the standard notion of regular cell-partition, also known as regular CW-complex, introduced by J.H.C. Whitehead (see e.g. Ziegler, 2002). Namely, a *regular cell-partition* is a finite semiregular cell-partition, where the closure of each cell is homeomorphic to a closed ball.

According to our definition, for example, the vertical projection on the plane of the graph (in \mathbb{R}^3) of a continuous piecewise-linear (more properly, piecewise-affine) function f on \mathbb{R}^2 , which is defined by finitely many affine equations and inequalities, naturally induces a semiregular cell-partition \mathcal{P}_f of \mathbb{R}^2 : each (closed) 2-cell is a maximal linearity set of the function. However, \mathcal{P}_f fails to be a regular cell-partition due to its unbounded cells. Semiregular partitions are especially well-suited for the study of the topology of real semialgebraic sets – any such set has a canonical finite semiregular cell-partition, known in this case as Whitney stratification.

It is easy to see that any finite semiregular cell-partition can be subdivided into a triangulation. It can be easily proven that a semi-regular cell-partition can also be subdivided

into a triangulation. However, since the algorithmic part of this paper deals only with finite partitions, we omit the proof.

A subset of \mathbb{R}^n is called polyhedral if it is defined by a propositional formula in the first-order (quantifiers can only be over numbers, not sets) language of the reals (\mathbb{R}) that uses only affine equations and inequalities. A subset S of $\mathbb{S}^n \subset \mathbb{R}^{n+1}$ is called polyhedral if $S = \mathbb{S}^n \cap E$, where $E \subset \mathbb{R}^{n+1}$ is defined by a propositional formula in the language of the reals (\mathbb{R}) that uses only linear (homogeneous) equations and inequalities.

Let $(\mathcal{M}, \mathcal{P})$ be a cell-partition of a manifold \mathcal{M} , and let $r : \mathcal{M} \rightarrow \mathbb{X}^n$ be continuous. The triple $(\mathcal{M}, \mathcal{P}; r)$ is called a *PL-realization* in \mathbb{X}^n if for each $C \in \mathcal{P}$ the set $r(\overline{C})$ is (1) polyhedral, and (2) homeomorphic to a closed k -ball ($k \leq \dim C$). We call such realization *dimension-preserving* if for each $C \in \mathcal{P}$ the image $r(\overline{C})$ is of affine dimension equal to $\dim C$ and r , restricted to \overline{C} , is a homeomorphism. Note that in this case although r need not even be an immersion, the restriction of r to the closure of each cell must be an embedding. $(\mathcal{M}, \mathcal{P}; r)$ is referred to as a *PL-hypersurface* if \mathcal{P} is a semiregular cell-partition and r is a dimension-preserving PL-realization. In particular, any dimension-preserving PL-realization in \mathbb{R}^n of the closure of a topological k -cell of $(\mathcal{M}, \mathcal{P})$ can always be implemented as a closed k -dimensional NEF-polyhedron homeomorphic to a k -ball. Thus, NEF-polyhedra provide a convenient basis for a computer implementation of PL-hypersurfaces in any dimension.

Since the terms *face* and *cell* are used both for abstract cells and their geometric realizations, when there is a need to stress the distinction we say *topological cell* or *geometric cell* respectively. Throughout the paper all geometric faces, just as all cells in topological partitions, are assumed to be *relatively open* (e.g. a 1-cell is an open segment, rather than a closed one) At times we refer to $(n - 1)$ -faces as *facets*, $(n - 2)$ -faces as *ridges*, and $(n - 3)$ -faces as *corners* (we also use these intuitive names for the topological cells of $(\mathcal{M}, \mathcal{P})$ – the intended meaning will be clear from the context).

From now on *all maps are continuous*. Here is our main theorem stated formally. Let \mathcal{M} ($\dim \mathcal{M} = n - 1 \geq 2$) be connected and let $(\mathcal{M}, \mathcal{P}; r)$ be a *PL-hypersurface*. Suppose $r : \mathcal{M} \rightarrow \mathbb{R}^n$ has at least one point of strict convexity; also, suppose that r is locally convex at all points of all corners of $(\mathcal{M}, \mathcal{P})$. Notice that if the last condition holds for some point of a corner, it holds for all points of the corner. Suppose also that each ball in \mathbb{R}^n contains finitely many faces of the surface. *Our main Theorem 13 states that under these conditions r is a homeomorphism onto the boundary of a convex body.* The idea of the proof of Theorem 13 is outlined at the end of Section 3. Theorem 13 implies a test for global convexity of a bounded closed PL-hypersurface that proceeds by checking the local convexity on each of the corners. The algorithm is given in Section 6. The complexity of our test depends not only on the model of computation, but also on the way the surface is given as input data. Let the input be the poset of facets, ridges, and corners. Suppose for each corner-ridge incidence (C, R) we are given a Euclidean inner normal to $r(R)$ at $r(C)$, and for each corner-facet incidence (C, F) we are given a Euclidean inner normal to $r(F)$ at $r(C)$. If we adopt the algebraic complexity model where each of scalar operations {sign determination, addition, subtraction, multiplication, and division} has unit cost, the complexity of the algorithm is

$O(nf_{n-3n-2}) = O(nf_{n-3n-1})$. Complexity under other models is discussed in Section 7, where we also indicate the cost of extracting the required input information from more common input representations.

3 Geometry of locally-convex immersions

Recall that a path joining points x and y in a topological space \mathcal{T} is a map $\alpha : [0, 1] \rightarrow \mathcal{T}$, where $\alpha(0) = x$ and $\alpha(1) = y$. Such a path is called an *arc* if α is injective. Denote by $\text{Arcs}_{\mathcal{M}}(x, y)$ the set of all arcs joining $x, y \in \mathcal{M}$.

Any realization $r : \mathcal{M} \rightarrow \mathbb{X}^n$ induces a distance d_r on \mathcal{M} by

$$d_r(x, y) = \inf_{\alpha \in \text{Arcs}_{\mathcal{M}}(x, y)} |r(\alpha)|,$$

where $|r(\alpha)| \in \mathbb{R}_+ \cup \infty$ stands for the length of the r -image of the arc α joining x and y on \mathcal{M} (we call it the r -distance, because it can take on the value ∞).

Of course, for a general realization r it is not clear *a priori* that there is a path of finite length on $r(\mathcal{M})$ joining $r(x)$ and $r(y)$ (where x and y are in the same connected component). Here we need to introduce a technical notion of *completeness*, which is essential to the correctness of van Heijenoort's theorem. $r : \mathcal{M} \rightarrow \mathbb{X}^n$ is called complete if any Cauchy sequence on \mathcal{M} (with respect to the r -distance) converges to a point of \mathcal{M} . Recall that a sequence $\{x_n\}$ is Cauchy under the distance d if for any real $\varepsilon > 0$ there is $N \in \mathbb{N}$ such that for any $n, m \geq N$ we have $d(x_n, x_m) \leq \varepsilon$. Completeness is a rather subtle notion: a space may be complete under a metric d and not complete under another metric d_1 , which is topologically equivalent to d (i.e. $x_n \xrightarrow{d} a$ iff $x_n \xrightarrow{d_1} a$).

Lemma 2 (*Van Heijenoort, 1952; pp. 227-228*) *Let $f : \mathcal{M} \rightarrow \mathbb{X}^n$ be a complete locally-convex immersion of an $(n - 1)$ -manifold \mathcal{M} . Then any two points in the same connected component of \mathcal{M} can be connected by an arc of finite length. The topology on \mathcal{M} defined by the f -distance is equivalent to the intrinsic (original) topology on \mathcal{M} .*

The input to the convexity checker need not be an immersed surface, since in the algorithm r is only expected to be a homeomorphism on the closures of cells, rather than on stars. Hence, the above lemma cannot be applied as stated. A realization is called *proper* if the preimage of every compact set is compact. A proper realization is always closed. For a given class of realizations it is usually much easier to check for properness than for completeness. Furthermore, the notion of properness is topological, while that of completeness is metrical. The following seems to be a folklore result (see e.g. Burago and Shefel, p. 50), so we do not give a proof here.

Lemma 3 *A proper realization of any manifold in \mathbb{X}^n is complete.*

Since we only consider maps r that send a finite number of cells into each bounded subset of \mathbb{R}^n , we do not have to worry about completeness. For locally-convex immersions the concepts of properness and completeness coincide:

Lemma 4 (Van Heijenoort) *A complete locally-convex immersion of a connected $(n - 1)$ -manifold into \mathbb{X}^n is proper.*

Van Heijenoort’s proofs of Lemmas 2 and 4 given in the original for \mathbb{R}^n are valid, word by word, for \mathbb{S}^n and \mathbb{H}^n , since these lemmas are entirely of local nature. If f is a locally-convex immersion, then for a “sufficiently small” subset \mathcal{S} of \mathcal{M} the map $f|_{\mathcal{S}}$ is a homeomorphism and, therefore, the topology on \mathcal{S} that is induced by the metric topology of \mathbb{X}^n is equivalent to the intrinsic topology of \mathcal{S} and, thanks to Lemma 2, to the f -distance topology. Thus, for sufficiently small subsets of \mathcal{M} (but not $i(\mathcal{M})$!) the three topologies considered in this section are equivalent – a fact that will be used throughout the text without an explicit reference to the above lemmas.

Theorem 5 (Van Heijenoort, 1952) *If a complete locally-convex immersion f of a connected $(n - 1)$ -manifold \mathcal{M} into \mathbb{R}^n ($n \geq 3$) has a point of strict convexity, then f is a homeomorphism onto the boundary of a convex body.*

Algorithmically, this means that when \mathcal{M} is known to be connected and r is known to be complete, to check convexity we have to check *all* points of \mathcal{M} for local convexity and, in addition, verify that at least one of them is a point of strict convexity. Let us now turn our attention to the world of PL-hypersurfaces, where confirming *local convexity at a point* is equivalent to confirming *local convexity at all points* of the (always unique on the level of $(\mathcal{M}, \mathcal{P})$) face containing this point. Denote by $B_R(p)$ a ball of radius $R > 0$ centered at a point $p \in \mathbb{X}^n$. It follows from the definition of dimension-preserving PL-realization that for any cell C of $(\mathcal{M}, \mathcal{P})$ the surface “looks the same” at all points of C : namely, for any $x, y \in C$ the intersections of $B_R(r(x))$ and $B_R(r(y))$ with $r(\text{Star } C)$ are congruent for sufficiently small R . Faces of dimensions $n - 1$ and $n - 2$ are the easiest to understand: if $\dim C = n - 1$, then $\text{Star } C = C$ and for any $x \in C$ we see that $B_R(r(x)) \cap r(\text{Star } C)$ is an $(n - 1)$ -ball in $\text{aff } r(C)$; if $\dim C = n - 2$, then $B_R(r(x)) \cap r(\text{Star } C)$ consists of two $(n - 1)$ -dimensional half-balls sharing a common $(n - 2)$ -ball (see Figure 9 for $n = 3$). Thus, any dimension-preserving PL-realization $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is locally-convex at all points of $(n - 1)$ - and $(n - 2)$ -cells; if $n = 3$ the only remaining cells are the 0-cells, i.e. the vertices. Thus, in the case of \mathbb{R}^3 we only have to check local convexity at the vertices. There is no need to check the existence of a point of strictly convexity in the compact case.

Lemma 6 *If $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is a locally-convex immersion of a compact connected $(n - 1)$ -manifold \mathcal{M} , then r has a point of strict convexity.*

Proof. As \mathcal{M} is compact and r is an immersion, $\text{conv } r(\mathcal{M})$ is a compact subset of \mathbb{R}^n . Since $\text{conv } r(\mathcal{M})$ is compact, it is also bounded and, in particular, does not contain lines. Any non-empty convex set, which is free of lines, has a non-empty set of extreme points (a point on the boundary of a convex set is extreme if it is not interior to any line segment contained

in set's boundary). Thus $\partial \operatorname{conv} r(\mathcal{M})$ contains an *extreme* point. Straszewicz's theorem (e.g. Rockafellar, 1997, p. 167) states that the *exposed* points of a closed convex set form a dense subset of *extreme* points of this set. Thus, $\operatorname{conv} r(\mathcal{M})$ has an exposed point. Since an exposed point y cannot be written as a *strict convex combination* of other points of the set, y must lie in $r(\mathcal{M})$. Let x be a point from $r^{-1}(y)$. Since r is locally-convex at x and there exists a hyperplane H through y that has empty intersection with $r(\mathcal{M}) \setminus y$, we conclude that the map r is strictly locally-convex at x . ■

In the case of PL-hypersurfaces in \mathbb{R}^n (where $n \geq 3$) it looks like we have to check the local convexity at all $(n - 3)$ -faces, $(n - 4)$ -faces, etc, all the way to the vertices. A trivial observation is that it is enough to check the local convexity at the vertices, since the local convexity at a vertex is sufficient for the local convexity at each of the faces at this vertex. However, we set to prove that for a bounded PL-hypersurface it is enough to check the local convexity only at the faces of dimension $n - 3$ – the rest will follow. We proceed as follows. First we notice (Section 4) that a locally-convex hypersurface in \mathbb{S}^n (where $n \geq 3$) satisfying the conditions of (Euclidean) Theorem 13 is a convex immersion. Using this result we reduce the general problem of local convexity testing to the local convexity testing at corners in Section 5 (recall that we have the local convexity at facets and ridges for free).

4 Locally-convex hypersurfaces in \mathbb{S}^n

Theorem 7 (below) of Jonker and Norman (1973) generalizes the one by Van Heijenoort by characterizing the case of non-convex locally-convex (complete and connected) immersions. Any such immersion is an immersion onto the product of a locally-convex, but not convex, plane curve and a complimentary affine subspace. We show that for locally-convex immersions into a sphere of dimension $n \geq 3$ the absence of points of strict convexity cannot result in the loss of global convexity, as it happens in the Euclidean case. The proof of our main theorem relies on the result of Jonker & Norman, although their theorem does not directly imply ours. One of the difficulties is that a compact convex set on the sphere may be free of extreme points!

Theorem 7 (*Jonker-Norman*) *Let $i : \mathcal{M} \rightarrow \mathbb{R}^n$ ($n \geq 3$) be a complete locally-convex immersion of a connected $(n - 1)$ -manifold. Then for any $x \in \mathcal{M}$ there is a unique submanifold \mathcal{D} through x such that*

- (1) $i(\mathcal{D}) = \operatorname{aff} i(\mathcal{D})$,
- (2) $i|_{\mathcal{D}}$ is a homeomorphism,
- (3) \mathcal{D} is maximal with respect to 1) and 2).

Furthermore, for any hyperplane H through x , which is complementary to $\operatorname{aff} \mathcal{D}$ the set $\mathcal{G} := i^{-1}(\mathcal{M} \cap H)$ is a submanifold of \mathcal{M} with $\dim \mathcal{M} = n - 1 - \dim \mathcal{D}$ such that

- (1) $\mathcal{M} = \mathcal{D} \times_{\operatorname{Top}} \mathcal{G}$,
- (2) $i|_{\mathcal{G}}$ is a locally-convex immersion into $\operatorname{aff}(\mathcal{M} \cap H)$ with at least one point of strict convexity,

(3) if \mathcal{D}' and \mathcal{G}' are to $x' \in \mathcal{M}$ what \mathcal{D} and \mathcal{G} are to x , then $\mathcal{D}' \cong_{\text{Top}} \mathcal{D}$, $\mathcal{G}' \cong_{\text{Top}} \mathcal{G}$, and $i(\mathcal{G})$ is equivalent to $i(\mathcal{G}')$ under the action of affine automorphisms of $i(\mathcal{M})$ that map \mathcal{D} to itself.

Finally, if i is not a convex embedding, then $\dim \mathcal{G} = 1$.

Let us look at Figure 1 (this is an embedding, so we do not have to distinguish between \mathcal{M} and $i(\mathcal{M})$): if $x = \mathbf{b}$, then $\mathcal{D} = \mathbf{l}$ and the 4-gone $(abcd)$ can be chosen as \mathcal{G} .

Whenever we have a map i that satisfies Jonker-Norman's theorem we will talk about the *locally-convex direct decomposition* of i ; we may also say that the immersion i *splits* into the *locally-convex direct product* of $i : \mathcal{D} \rightarrow D$ and $i : \mathcal{G} \rightarrow G$. When $i(\mathcal{G})$ is chosen to be perpendicular to $i(\mathcal{D})$, we call the decomposition *orthogonal*.

Assume that \mathbb{S}^n is embedded as the standard sphere into \mathbb{R}^{n+1} .

Theorem 8 *Let $i : \mathcal{M} \rightarrow \mathbb{S}^n$ ($n \geq 3$) be a complete locally-convex immersion of a connected $(n - 1)$ -manifold \mathcal{M} . Then $i(\mathcal{M}) = \mathbb{S}^n \cap \partial K$, where K is a convex cone in \mathbb{R}^{n+1} having the origin as its apex.*

Note that the statement of Theorem 8 is invalid for $n = 2$. For example, although the 1-surface in \mathbb{S}^2 depicted in Figure 5 is locally-convex at all points, it does not bound any convex set on \mathbb{S}^2 . The rest of Section 4 is somewhat technical and some readers may want to skip it.




Fig. 5. The four bold arcs form a non-convex, but locally-convex polygon in \mathbb{S}^2

4.1 Proof of the local-to-global convexity theorem for \mathbb{S}^n

While discussing immersions, it is important to remember that they need not be injective; for example, we do not really consider, say, a line L on the surface defined by i as a subset of $i(\mathcal{M})$, but rather the map $i : \mathcal{L} \rightarrow L$, where \mathcal{L} is a 1-submanifold of \mathcal{M} and $L = i(\mathcal{L})$. The same philosophy applies to any geometric subobject of $i : \mathcal{M} \rightarrow \mathbb{X}^n$. In the case of points we use the shorthand $i(x)$ to mean $i : x \rightarrow i(x)$. On occasions, when for some $\mathbf{p} \in i(\mathcal{M})$ it is absolutely clear from the context as to which point x of $i^{-1}\mathbf{p}$ is considered, we refer to $i : x \rightarrow \mathbf{p}$ simply as “point \mathbf{p} ”.

By analogy with the traditional terminology for ruled surfaces and cylinders in 3D we refer to $i : \mathcal{D} \rightarrow D$, where $D = i(\mathcal{D})$, as a *directrix* and to $\vec{D} = D - D$ as the *linear directrix*

of $i : \mathcal{M} \rightarrow \mathbb{X}^n$. Similarly, we refer to $i : \mathcal{G} \rightarrow G$, where $G = i(\mathcal{G})$, as a *generatrix* of $i : \mathcal{M} \rightarrow \mathbb{X}^n$.

By the Jonker-Norman theorem \mathcal{D} can be chosen so that it passes through a point of strict convexity of $i|_{\mathcal{G}}$, in which case we call it an *exposed directrix*. Note that such a directrix is *never interior* to any flat of *higher dimension* contained in $i : \mathcal{M} \rightarrow M = i(\mathcal{M})$. When needed we refer to D as the geometric directrix and to \mathcal{D} as an abstract directrix, etc. In Figure 1 the line \mathbf{l} is a directrix and the 4-gone $(abcd)$ is a generatrix. We will need the following corollary of Theorem 7.

We call a connected submanifold \mathcal{S} of a topological space \mathcal{T} *flat* with respect to a realization map $r : \mathcal{T} \rightarrow \mathbb{X}^n$ if $r : \mathcal{S} \rightarrow r(\mathcal{S})$ is a homeomorphism onto a subspace of \mathbb{X}^n of the same dimension as \mathcal{S} ; in this case we call $r(\mathcal{S})$ a *flat* contained in (\mathcal{T}, r) .

Corollary 9 *In the context of Jonker-Norman theorem any flat containing an exposed point of $i|_{\mathcal{G}}$ is contained in the exposed directrix through this point.*

The spherical convexity criterion, Theorem 8, is a direct consequence of Theorem 10 and Theorem 12; the former deals with the case where a point of strict convexity is absent and the latter deal with the case where it exists. The idea of proof of Theorem 10 is to apply Jonker-Norman's theorem locally, i.e. for a finite number of open hemispheres covering \mathbb{S}^n . The hypersurface, considered over each such hemisphere has a number (possibly 0) of connected components, each of which having a unique orthogonal Jonker-Norman decomposition (since the affine geometry of a hemisphere is essentially equivalent to the geometry of \mathbb{R}^n). Among all such connected pieces of (\mathcal{M}, i) lying in different hemispheres we pick one that has the exposed directrix of minimal dimension. The Jonker-Norman decomposition is so "rigid" that whenever an exposed directrix continues from a hemisphere H to a hemisphere H' ($H \cap H' \neq \emptyset$), the Jonker-Norman decompositions on $H \cap H'$ that are inherited from H and H' must agree. As a result we get an analog of Jonker-Norman's theorem for the sphere. Theorem 12 is proven by a combination of topological considerations and a metric (perturbation type) argument reducing the problem to the Euclidean one.

Theorem 10 *Let $i : \mathcal{M} \rightarrow \mathbb{S}^n$ ($n \geq 3$) be a complete locally-convex immersion of a connected $(n - 1)$ -manifold \mathcal{M} without points of strict convexity. Then $i(\mathcal{M}) = \mathbb{S}^n \cap \partial K$, where K is a convex cone in \mathbb{R}^{n+1} having the origin as its apex.*

For a hemisphere $H \subset \mathbb{S}^n \subset \mathbb{R}^{n+1}$ we denote by c_H the central spherical projection map (see Figure 6) from H onto the tangent n -plane \mathbf{T}_H to \mathbb{S}^n at the center of H (when we find it convenient to index the space and the projection map by the center \mathbf{p} of H we write $\mathbf{T}_{\mathbf{p}}$ and $c_{\mathbf{p}}$ instead of \mathbf{T}_H and c_H).

\mathbf{T}_H is an *affine* real n -space; when we need to treat it as a *linear* space (i.e. $\overrightarrow{\mathbf{T}}_H = \mathbf{T}_H - \mathbf{T}_H$), we identify the origin of $\overrightarrow{\mathbf{T}}_H$ with the center of H . First, let us make the following trivial observation.




Fig. 6. Central spherical projection map from hemisphere H to \mathbf{T}_H .

Lemma 11 *Let \mathbf{p} be an extreme point of a convex set $B \subset \mathbb{S}^n$. Then every neighborhood of \mathbf{p} has an exposed point of B . In particular, if \mathbf{p} is not exposed itself, there are infinitely many exposed points on ∂B arbitrarily close to \mathbf{p} .*

Proof. This lemma is essentially a restatement for spherical spaces of a well-known theorem of Straszewicz on convex sets in \mathbb{R}^n (Rockafellar, p.167). Since our lemma is entirely of local nature, the proof in Rockafellar (1997) applies without changes. Alternatively, consider the tangent space $\mathbf{T}_{\mathbf{p}}$ to $\mathbb{S}^n \subset \mathbb{R}^{n+1}$ at \mathbf{p} . The central projection maps the hemisphere H centered at \mathbf{p} onto $\mathbf{T}_{\mathbf{p}}$ and $H \cap B$ onto a convex set $B_{\mathbf{p}}$ in $\mathbf{T}_{\mathbf{p}}$. The Euclidean theorem can now be applied to \mathbf{p} as an extreme point of the convex set $B_{\mathbf{p}}$ in $\mathbf{T}_{\mathbf{p}} \cong \mathbb{R}^n$. The property of a point to be extreme or exposed is preserved under the central projection and its inverse. ■

Proof of Theorem 10. If there is $x \in \mathcal{M}$ such that $i(x)$ is extreme for some convex witness at x , then, by Lemma 11, there is an exposed point at every neighborhood of x . Since an exposed point of a convex body is the same as a point of strict convexity of its surface, (\mathcal{M}, i) is strictly convex in at least one point, which is impossible. Thus, for each $x \in \mathcal{M}$ the image $i(x)$ is an interior point of a segment on $M := i(\mathcal{M})$.

For an open hemisphere H let $\mathcal{S}H$ denote the set of all maximal connected submanifolds of $i^{-1}(H)$ whose i -images lie in H . When $\mathcal{S}H \neq \emptyset$ each $\mathcal{S} \in \mathcal{S}H$ is an $(n-1)$ -submanifold of \mathcal{M} and $c_H \circ i : \mathcal{S} \rightarrow \mathbf{T}_H$ is a complete locally-convex immersion (e.g. by Lemmas 4 and 3).

So, for any such $\mathcal{S} \in \mathcal{S}H$ by Jonker-Norman's theorem the map $c_H \circ i : \mathcal{S} \rightarrow \mathbf{T}_H$ has a *locally-convex direct decomposition* $\mathcal{S} = \mathcal{G} \times \mathcal{L}$, where $c_H \circ i|_{\mathcal{G}}$ is a locally-convex immersion of a compact connected g -submanifold \mathcal{G} and $c_H \circ i : \mathcal{L} \rightarrow L$ is a homeomorphism from a d -submanifold $\mathcal{L} \subset \mathcal{S}$ onto a d -subspace of \mathbf{T}_H ($n-1 = d+g$). Denote by $\overrightarrow{L}(\mathcal{S})$ the linear space $L-L$.

Let us pick \mathcal{G} so that $G \perp L$ where $G := c_H \circ i(\mathcal{G})$. Then $c_H \circ i|_{\mathcal{S}}$ is the *orthogonal direct product* of the generatrix $c_H \circ i|_{\mathcal{G}}$ and the directrix $c_H \circ i|_{\mathcal{D}}$. On the hemisphere H this decomposition corresponds to the orthogonal locally-convex split of $i|_{\mathcal{S}}$ into a hemispherical generatrix $i : \mathcal{G} \rightarrow c_H^{-1}G$ and a hemispherical directrix $i : \mathcal{D} \rightarrow c_H^{-1}L$.

Let \mathfrak{H} be a finite covering of \mathbb{S}^n by *open hemispheres*. Since i does not have points of strict convexity, $i(\mathcal{M})$ is not completely covered by any single hemisphere. We will use \mathcal{S}_H for an element of $\mathcal{S}H$ – subindex H only indicates that \mathcal{S}_H was chosen from $\mathcal{S}H$. Likewise, once \mathcal{S}_H

is fixed, we may use $L_{\mathcal{S}_H}$ and $\mathcal{G}_{\mathcal{S}_H}$, etc. to indicate that $L_{\mathcal{S}_H}$ and $\mathcal{G}_{\mathcal{S}_H}$ are obtained from the direct decomposition of $c_H \circ i|_{\mathcal{S}_H}$.

Suppose $U = i(\mathcal{U})$ is a convex hypersurface for some connected submanifold $\mathcal{U} \subset \mathcal{S}_H$. Let $H, H' \in \mathfrak{H}$, and $H \cap H' \cap U \neq \emptyset$. Then there is a unique $\mathcal{S}_{H'} \in \mathcal{S}H'$ such that $\mathcal{S}_{H'}$ contains $i^{-1}H' \cap \mathcal{U}$. We will refer to this fact by saying that *whenever a convex subsurface of $i : \mathcal{S}_H \rightarrow H$ protrudes into H' , the surface $i : \mathcal{S}_H \rightarrow H$ extends uniquely into $H' \cup H$ along U* (or, in other words, the map $i|_{\mathcal{S}_H}$ extends uniquely over $i^{-1}H'$ along $\mathcal{U} \cap i^{-1}H'$). In this context $\mathcal{S}_{H'} \cup \mathcal{S}_H$ is called the *extension* of \mathcal{S}_H and $\mathcal{S}_{H'}$ is called an *adjoint* of \mathcal{S}_H .

Among the elements of \mathfrak{H} that overlap with $i(\mathcal{M})$, let H_0 be one where we can pick $\mathcal{S}_{H_0} \in \mathcal{S}H_0$ so that $d := \dim \vec{L}(\mathcal{S}_{H_0}) \leq \dim \vec{L}(\mathcal{S}_H)$ for all H that overlap with $i(\mathcal{M})$ and each possible choice of $\mathcal{S}_H \in \mathcal{S}H$; let $i : \mathcal{D}_0 \rightarrow D_0 = i(\mathcal{D}_0)$ be an exposed hemispherical directrix for \mathcal{S}_{H_0} .

If $d = 0$, then $\mathcal{S} = \mathcal{G}$, where $c_H \circ i : \mathcal{G} \rightarrow \mathbf{T}_H$ has a point of strict convexity, which contradicts our assumption about i .

If $d = n - 1$, then $\text{aff } D_0$ is an $(n - 1)$ -hemisphere of H_0 and $\mathcal{S}_{H_0} = \mathcal{D}_0$. (A k -hemisphere in \mathbb{S}^n is a hemisphere of a k -dimensional subspace of \mathbb{S}^n .) We know that whenever a convex subsurface of $i : \mathcal{S}_{H_0} \rightarrow H_0$ protrudes into H , the surface $i : \mathcal{S}_{H_0} \rightarrow H_0$ extends uniquely into H along this subsurface, which implies that $i|_{\mathcal{D}_0}$ can be extended to all hemispheres overlapping with $\text{aff } D_0$. Since \mathcal{M} is a connected $(n - 1)$ -manifold, $i : \mathcal{M} \rightarrow \mathbb{S}^n$ is an immersion onto $\text{aff } \mathcal{D}_0$, which is, by the covering mapping (see Seifert & Threlfall) theorem, a homeomorphism if $n > 2$.

Let $1 \leq d \leq n - 2$. Let $H \cap D_0 \neq \emptyset$ for some $H \in \mathfrak{H}$. We know $i|_{\mathcal{S}_{H_0}}$ extends in a unique way along $D_0 \cap H$ into $H_0 \cup H$: denote the adjoint element of $\mathcal{S}H$ by \mathcal{S}_H . Obviously, $\text{aff } D_0 \cap H$ is an extreme (geometric) hemispherical directrix for $i|_{\mathcal{S}_H}$. As D_0 is completely covered by elements of \mathfrak{H} , the submanifold \mathcal{D}_0 extends to a connected component of $i^{-1}(\text{aff } D_0)$ inside of \mathcal{M} . Set $D := \text{aff } D_0$ and let \mathcal{D} be a maximal connected d -submanifold of \mathcal{M} such that $D = i(\mathcal{D})$. Since i is a complete immersion, it is proper (preimages of compact sets are compact) and \mathcal{D} is compact. Thus, the preimage of any $\mathbf{p} \in D$ under $i|_{\mathcal{D}}$ is a finite set of size that does not depend on \mathbf{p} .

Without loss of generality we assume that D is completely covered by hemispheres $H_0, \dots, H_N \in \mathfrak{H}$, all centered at points of D ; denote this subset of \mathfrak{H} by \mathfrak{H}_D . Let \mathcal{S} be a connected component of $i^{-1}(\cup_{j=0}^N H_j)$ that contains \mathcal{D} , i.e. \mathcal{S} is the unique maximal extension of \mathcal{S}_{H_0} into $\cup_{j=0}^N H_j$ along D . For $H_k, H_l \in \mathfrak{H}_D$, where $H_k \cap H_l \neq \emptyset$, on each connected component of $i^{-1}(H_k \cap H_l) \cap \mathcal{S}$ the locally-convex orthogonal decompositions of $i|_{\mathcal{S}}$, which are induced by the restrictions of i to $\mathcal{S} \cap i^{-1}H_k$ and $\mathcal{S} \cap i^{-1}H_l$ respectively, agree; this follows directly from Jonker-Norman's theorem. Furthermore, since both H_k and H_l are centered at D , the generatrices in these two locally-convex *orthogonal* decompositions are all isometric to each other – the rotational subgroup $\text{Iso}^+(D)$ of $\text{Iso}(D)$ is transitive on them.

Thus, we have a *locally-convex orthogonal fibration* of the immersion $i|_{\mathcal{S}}$: namely, we have

a continuous map $\pi : \mathcal{S} \rightarrow \mathcal{D}$, which sends each (topological) generatrix into its base point on \mathcal{D} and for each $x \in \mathcal{D}$ there is a neighborhood $\mathcal{U}_x \subset \mathcal{D}$ such that $\pi^{-1}(\mathcal{U}_x)$ is the direct orthogonal product of \mathcal{U}_x and a fiber \mathcal{G}_x over x , such that $i|_{\mathcal{G}_x}$ is a locally-convex immersion into D_x^\perp , where the latter is the orthogonal complement of D through x . Inside of each $H \in \mathfrak{H}_D$ the fibers (i.e. generatrices) are isometric; moreover, as we just noticed above, the fibers from different H 's are also isometric. Thus, the constructed locally-convex fibration of the immersion $i|_{\mathcal{S}}$ is, in fact, a direct product decomposition, i.e. $\mathcal{S} = \mathcal{D} \times \mathcal{G}$, where $\dim \mathcal{G} = n - 1 - d$ and $i|_{\mathcal{G}}$ is a locally convex immersion into a $(n - d)$ -hemisphere perpendicular to D .

Set $D^* := \mathbb{S}^n \cap (\text{cone } D)^\perp$, where $\text{cone } D$ is the cone with apex at $\mathbf{0}$ over D . D^* consist of all points of \mathbb{S}^n that are not covered by the elements of \mathfrak{H}_D . We claim that all generatrices from the orthogonal decomposition of $i|_{\mathcal{S}}$ “reach to D^* ”, i.e. for each generatrix $\mathcal{G} \subset \mathcal{S}$ and any neighborhood of D^* there is $p \in \mathcal{G}$ such that $i(p)$ lies in this neighborhood. By contradiction: let $p \in \mathcal{G} \subset \mathcal{S}$ be such that the distance $\rho > 0$ between $i(p)$ and D^* is equal to the distance between $i(\mathcal{G})$ and D^* . Since all generatrices are isometric with respect to $Iso^+(D)$, \mathcal{S} contains a submanifold mapped onto the orbit of p under the induced action of $Iso^+(D)$ on \mathcal{S} , but does not contain any points mapped by i to spherical points at the distance smaller than ρ from D^* . Then i is not locally-convex at all points of this submanifold. Thus, all generatrices of the orthogonal decomposition of $i|_{\mathcal{S}}$ “reach to D^* ”.

Since \mathcal{M} is compact, for any $p \in \bar{\mathcal{S}} \setminus \mathcal{S}$ we have $i(p) \in D^* \cong \mathbb{S}^{n-d-1}$. Then $i(p)$ belongs to the closure of each generatrix from the orthogonal fibration of $i|_{\mathcal{S}}$ with base $i|_{\mathcal{D}}$. Let $H(\mathbf{p})$ be the hemisphere centered at $\mathbf{p} = i(p)$. Under $c_H(\mathbf{p}) : H_p \rightarrow \mathbf{T}_{H(\mathbf{p})}$ the points of D correspond to rays emanating from the origin of $\mathbf{T}_{H(\mathbf{p})}$, or, in other words, in “the world of” \mathbf{T}_p the spherical subspace D corresponds to a “ d -sphere at infinity”, which we denote by $D(\mathbf{T}_{H(\mathbf{p})})$. Thus, the isometry group of the surface $(\bar{\mathcal{S}}, c_H(\mathbf{p}) \circ i)$ includes all linear isometries (in particular, rotations about \mathbf{p}) that preserve the sphere $D(\mathbf{T}_{H(\mathbf{p})})$ at infinity. We know that $\mathbf{p} = i(p) = c_H(\mathbf{p}) \circ i(p)$ must belong to the interior of a segment $I = c_H(\mathbf{p}) \circ i(J)$ on this surface. Any isometry of $(\bar{\mathcal{S}}, c_H(\mathbf{p}) \circ i)$ will map I to another line segment. Because of local convexity at p , the isometries preserving the sphere $D(\mathbf{T}_{H(\mathbf{p})})$ at infinity belong to the isometry group of a supporting hyperplane at $\mathbf{p} = c_H(\mathbf{p}) \circ i(p)$. Since I must be in this hyperplane, $i(p)$ is interior to a $(d + 1)$ -flat of $\bar{\mathcal{S}}$. We will have to deal separately with the cases $d = n - 2$ and $d \leq n - 3$.

Case: $d = n - 2$. $\bar{\mathcal{S}} \setminus \mathcal{S} \cong \mathbb{S}^0$. Then $i(p)$ is interior to an $(n - 1)$ -flat $i : \mathcal{F} \rightarrow F$ of $(\bar{\mathcal{S}}, i)$. We will show that $i|_{H(\mathbf{p})}$ is a homeomorphism onto an open $(n - 1)$ -hemisphere of \mathbb{S}^n . Consider a directed geodesic in $\text{aff } F$ with source at $i(p)$ that does not extend to D inside of $(\bar{\mathcal{S}}, i)$. Let $\mathbf{b} = i(b)$ be a point where this geodesic first diverges with the surface $(\bar{\mathcal{S}}, i)$. Point b belongs to a unique fiber from the orthogonal fibration of $(\bar{\mathcal{S}}, i)$ with base (\mathcal{D}, i) . The isometry group of D is transitive on the fibers. Thus, all directed geodesics through $i(p)$ that lie in $(\bar{\mathcal{S}}, i)$ diverge with the $(n - 1)$ -flat F at the same distance from $i(p)$; hence, \mathcal{F} is an $(n - 1)$ -ball (in the i -distance) centered at p . But then all points of $\partial\mathcal{F}$ are extreme points for $(\bar{\mathcal{S}}, i)$, which contradicts our assumptions. Thus F is an $(n - 1)$ -hemisphere centered at $i(p)$ and

bounded by D . The same argument is applied to the other point of $\bar{\mathcal{S}} \setminus \mathcal{S} \cong \mathbb{S}^0$. Thus, i is a homeomorphism onto the surface made of two $(n - 1)$ -hemispheres glued together at their common $(n - 2)$ -dimensional boundary D .

Case: $1 \leq d < n - 2$. The central projection of a generatrix $i|_{\mathcal{G}}$ (where $G = i(\mathcal{G})$ and $G \perp D$) with a base point $i(x) = \mathbf{x} \in D$ onto its tangent subspace $\mathbf{T}_G \subset \mathbf{T}_{\mathbf{x}}$ at $\mathbf{x} \in D$ is a locally-convex unbounded complete surface $c_{\mathbf{x}} \circ i|_{\mathcal{G}}$. Since the topological dimension of the generatrix is larger than one, by Jonker-Norman's theorem it is an embedded convex surface in \mathbf{T}_G and \mathbf{x} is a (geometric) point of strict convexity for this surface. Thus, $i|_{\mathcal{G}}$ is a convex surface on $H_{\mathbf{x}}$. We need to understand the geometry of $i|_{\mathcal{G}}$ at infinity, i.e. at $\partial H_{\mathbf{x}} \cap D^*$. Because of strict convexity at $\mathbf{x} = i(x)$, $\partial H_{\mathbf{x}} \cap i(\bar{\mathcal{G}})$ is the boundary of a strictly convex compact set in $\partial H_{\mathbf{x}}$. Suppose $z \in \mathcal{G}$ is a point of strict convexity of $i|_{\mathcal{G}}$ and $z \neq x$. Then there is an exposed hemispherical directrix $i : \mathcal{D}_z \rightarrow H_p$ through z , distinct from $\mathcal{D} \cap i^{-1}H_{\mathbf{x}}$. Since $i(\mathcal{D})$ and $i(\mathcal{D}_z)$ are parallel in $H_{\mathbf{x}}$, they “intersect at infinity” (i.e. on $\partial H_{\mathbf{x}}$) over a common $(d - 1)$ -sphere. Thus, we have *two distinct exposed directrices* through the same point of \mathcal{M} . This is impossible by Corollary 9. Thus, $i|_{\mathcal{G}}$ has a unique point of strict convexity.

$c_{\mathbf{x}} \circ i : \mathcal{G} \rightarrow \mathbf{T}_G$ is an embedded unbounded complete convex hypersurface with a unique point of strict convexity. Let us apply a projective transformation that sends \mathbf{T}_G to another subspace \mathbf{P} of the same dimension in \mathbb{R}^{n+1} in such a way that the point \mathbf{x} of $G \subset \mathbf{T}_G$ is mapped to a point at infinity of \mathbf{P} . This will give us an embedded unbounded complete convex hypersurface in \mathbf{P} without points of strict convexity. By Jonker-Norman this hypersurface in \mathbf{P} is the product of a line L in \mathbf{P} and a compact convex hypersurface in a subspace of \mathbf{P} , which is complimentary to L . Thus, $c_{\mathbf{x}} \circ i(\mathcal{G})$ is the boundary of a cone with apex at \mathbf{x} over a convex compact set “on the sphere at infinity of \mathbf{T}_G ”. Hence, $c_{\mathbf{x}} \circ i|_{\bar{\mathcal{G}}}$ is an immersion onto a cone over a convex compact surface of topological dimension $(n - 1) - d - 1 = n - d - 2$ on D^* . Since all generatrices are isometric to $i|_{\bar{\mathcal{G}}}$ with respect to the action of $Iso^+(D)$, we conclude that \mathcal{M} contains a closed $(n - 1)$ -submanifold $\bar{\mathcal{S}}$ (without boundary). Since \mathcal{M} is connected, $\bar{\mathcal{S}} = \mathcal{M}$. ■

Remark on injectivity. The proof does not imply that i is an embedding. When there are no points of strict convexity, non-injectivity is possible if and only if $d = \dim \mathcal{D} = 1$. When $d > 1$ the classical covering mapping theorem (see e.g. Seifert-Threlfall) implies that the map is one-to-one.

Theorem 12 *Let \mathcal{M} be connected and let $i : \mathcal{M} \rightarrow \mathbb{X}^n$ be complete, locally-convex and, also, strictly locally-convex at $o \in \mathcal{M}$. Then $i : \mathcal{M} \rightarrow \mathbb{X}^n$ is a convex embedding.*

Proof. If $\mathbf{H}_o \subset \mathbb{S}^n$ is a supporting hyperplane at $i(o)$, then let us denote the *open* hemisphere defined by \mathbf{H}_o that contains the image of a small neighborhood of o by \mathbf{H}_o^+ ; the other open hemisphere is then denoted by \mathbf{H}_o^- . If \mathcal{N} is a neighborhood of x we denote by \mathcal{N} its punctured version, i.e. $\mathcal{N} \setminus x$.

Let \mathcal{S} be a maximal connected $(n - 1)$ -submanifold of \mathcal{M} such that $o \in \mathcal{S}$ and $i(\dot{\mathcal{S}}) \subset \mathbf{H}_o^+$.

Suppose that there is no $x \in \bar{\mathcal{S}} \setminus o$ with $i(x) \in \mathbf{H}_o$. Then the distance between $\bar{\mathcal{S}} \setminus \mathcal{N}_o$ (where \mathcal{N}_o is a small neighborhood of o) and \mathbf{H}_o is strictly positive. This means we can perturb \mathbf{H}_o so that $i(\bar{\mathcal{C}}_o)$ is in $\widetilde{\mathbf{H}}^+$, where $\widetilde{\mathbf{H}}$ is a perturbed version of \mathbf{H}_o . Let c be the central projection on $\mathbf{T}_{\widetilde{\mathbf{H}}^+}$ from $\widetilde{\mathbf{H}}^+$. Clearly, $c \circ i|_{\mathcal{S}}$ satisfies the conditions of Van Heijenoort's theorem; hence, $c \circ i|_{\mathcal{S}}$ is a convex embedding. Since \mathcal{M} is connected, $\mathcal{S} = \mathcal{M}$.

When a minimal geodesic between \mathbf{p} and \mathbf{q} is unique, we denote it by $[\mathbf{p}, \mathbf{q}]$; we will also use $[p, q]$, where $i(p) = \mathbf{p}$, $i(q) = \mathbf{q}$, to refer to a curve in \mathcal{M} that is mapped homeomorphically onto $[\mathbf{p}, \mathbf{q}]$. Let now $p \in \bar{\mathcal{S}} \setminus o$ be such that $i(p) \in \mathbf{H}_o$. If $i(p) \neq i(o)^{\text{op}}$, the opposite of $i(o)$, then the minimal geodesic joining $i(o)$ and $i(p)$ is unique and lies in \mathbf{H}_o . Let $\{i : [o, x_m] \rightarrow [i(o), i(x_m)]\}_{m \in \mathbb{N}}$, with $[o, x_m] \subset \mathcal{M}$, be a sequence of minimal geodesics that converges to $i : [o, p] \rightarrow [i(o), i(p)]$. The geodesics in this sequence lie arbitrarily close to \mathbf{H}_o . Since (\mathcal{M}, i) is strictly convex at o , we find that $p = o$, which contradicts the choice of p . Thus, the points of $\bar{\mathcal{S}} \setminus o$ that are mapped to \mathbf{H}_o are mapped to $i(o)^{\text{op}}$. Since i is a proper immersion, the preimage of $i(o)^{\text{op}}$ in \mathcal{M} is finite. Hence, the preimage of $i(o)^{\text{op}}$ in $\partial\mathcal{S} = \bar{\mathcal{S}} \setminus \mathcal{S}$ is finite. Clearly, $c \circ i|_{\mathcal{S}}$ satisfies the conditions of Jonker-Norman's theorem. Since i is strictly convex at o , $c \circ i|_{\mathcal{S}}$ must be a convex unbounded *embedding* onto a cylinder in $\mathbf{T}_{\mathbf{H}_o^+}$. The directrix must be 1-dimensional, for the cylinder has only two points at infinity, $i(o)$ and $i(o)^{\text{op}}$ (see Fig. 7). Thus, \mathcal{S} contains a punctured neighborhood of p homeomorphic to an $(n-1)$ -ball. If we add p to \mathcal{S} we get a compact connected $(n-1)$ -submanifold of \mathcal{M} (without boundary). Since \mathcal{M} is connected, $\bar{\mathcal{S}} = \mathcal{M}$. Thus, $i : \mathcal{M} \rightarrow \mathbb{X}^n$ is a convex embedding. ■

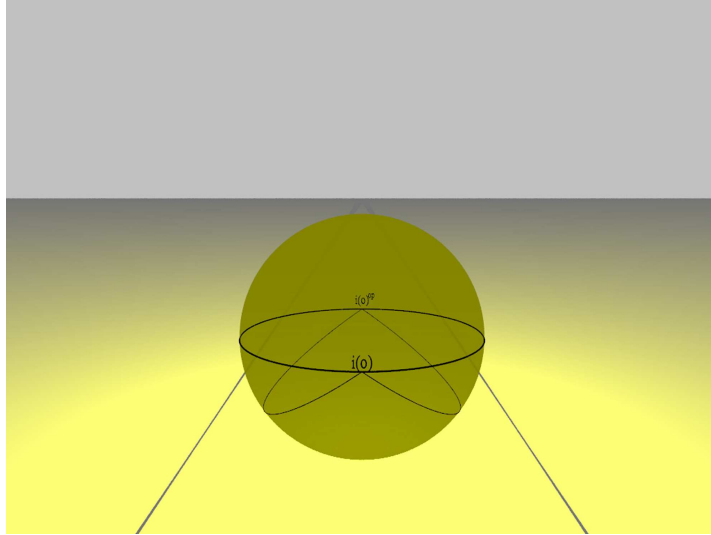


Fig. 7. The surface $i|_{\bar{\mathcal{S}}}$ has only two points on \mathbf{H}_o : $i(o)$ and $i(o)^{\text{op}}$

5 Locally-convex PL-surfaces in \mathbb{R}^n

Let \mathcal{P} be a fixed star-finite semi-regular cell-partition of \mathcal{M} . Recall that in our terminology a cell is always homeomorphic to an *open ball*. We say that r is locally-convex at a cell $C \in \mathcal{P}$ if it is locally-convex at each point of C .

Theorem 13 (main) *Let $r : \mathcal{M} \rightarrow \mathbb{R}^n$ ($n \geq 3$) be a dimension-preserving PL-realization of a connected manifold \mathcal{M} ($\dim \mathcal{M} = n - 1$) such that*

- (1) *r is locally-convex on each $(n - 3)$ -cell,*
- (2) *$r(\mathcal{M})$ is bounded or r is strictly locally-convex in at least one point of \mathcal{M} ,*
- (3) *r is proper.*

Then $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is an embedding onto the boundary of a convex body defined by (possibly infinitely many) affine inequalities.

Proof. Since r is proper, it is complete. We know that $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is locally-convex at all $(n - 3)$ -cells. Since r is dimension-preserving, it is locally-convex at all $(n - 1)$ - and $(n - 2)$ -cells; note that *any* dimension-preserving PL-realization of a manifold is *convex* at all points of all $(n - 1)$ - and $(n - 2)$ -cells.

Note that if r is bounded, then, by Lemma 6, the map r has a point of strict convexity. Thus, if we prove that r is locally-convex at all cells, by Theorem 5 the map $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is a convex embedding. We proceed by reverse induction in cell's dimension. Let $0 \leq k \leq n - 3$ and suppose that we have shown $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is locally-convex at all cells of dimension k and up. If $n - 3 = 0$, the proof is finished. So, let $n \geq 4$ and let us consider a $(k - 1)$ -cell $F \in \mathcal{P}$. Since \mathcal{P} is star-finite, $\text{Star } F$ contains finitely many cells. Consider $r(\text{Star } F) \cap \mathbb{S}_F$, where \mathbb{S}_F is a sufficiently small $(n - k)$ -sphere lying in a subspace complementary to $\text{aff } r(F)$ and centered at some point of $r(F)$. Note that $\dim \mathbb{S}_F = n - k \geq 2$. The surface $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is locally-convex at F if and only if the hypersurface $r : \mathcal{S} \rightarrow \mathbb{S}_F$, where $r(\mathcal{S}) = \mathbb{S}_F \cap r(\text{Star } F)$ and \mathcal{S} is the connected component of $r^{-1}(\mathbb{S}_F \cap r(\text{Star } F))$ that is contained in $\text{Star } F$, is a convex *embedding*.

Since $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is locally-convex at each k -cell, the surface $r : \mathcal{S} \rightarrow \mathbb{S}_F$ is locally-convex at the vertices, which correspond to the intersections of k -cells of (\mathcal{M}, r) with \mathbb{S}_F . Thus $r : \mathcal{S} \rightarrow \mathbb{S}_F$ is locally-convex. The completeness condition is clearly respected. By Theorem 8 $r : \mathcal{S} \rightarrow \mathbb{S}_F$ is a convex *immersion*. Notice that the condition $\dim \mathbb{S}_F = n - k > 2$ is essential to the applicability of Theorem 8 (see Figure 5 for a locally-convex surface in \mathbb{S}^2 which is not a convex surface in \mathbb{S}^2). If we can show that this immersion is an embedding, we have the local convexity at F . Recall that $n \geq 3$ and, therefore, \mathcal{M} and $\mathcal{M} \setminus \text{Sk}_{n-3}(\mathcal{M})$ have isomorphic fundamental groups. Thus, $r|_{\mathcal{M} \setminus \text{Sk}_{n-3}(\mathcal{M})}$ is a covering mapping. By the covering mapping theorem (see e.g. Seifert & Threlfall, 1980) it must be a homeomorphism. Thus, $r : \mathcal{S} \rightarrow \mathbb{S}_F$ is a bijection and the local convexity of r at F is proven. We proved that local convexity at k -cells implies local convexity at $(k - 1)$ -cells. Hence, r is locally-convex everywhere. Upon applying Van Heijenoort's Theorem 5 we conclude that r is a convex embedding. ■

Corollary 14 *Let $r : \mathcal{M} \rightarrow \mathbb{R}^n$ ($n > 2$) be a complete dimension-preserving PL-realization of a connected $(n - 1)$ -manifold \mathcal{M} , which is locally-convex at all $(n - 3)$ -cells. Suppose r is bounded or is strictly locally-convex in at least one point. If $(\mathcal{M}, \mathcal{P})$ is $(n - 3)$ -primitive, i.e. exactly $\mathfrak{3}$ $(n - 1)$ -cells make contact at each $(n - 3)$ -cell, then $r(\mathcal{M})$ is the boundary of a convex polyhedron.*

In particular, if we can realize $(\mathbb{S}^{n-1}, \mathcal{P})$, where \mathcal{P} is a regular $(n-3)$ -primitive cell-partition, in \mathbb{R}^n so that each k -cell is embedded as a set of affine dimension k for all k , then either this realization is a convex polytope or it is projectively equivalent to another realization which is a convex polytope (for the latter see Rybnikov, Zaslavsky, 2005). It would be interesting to apply this observation to open problems about convex 4-polytopes. Convex 4-polytopes are not well-understood, unlike their 3-dimensional counterparts, which are completely characterized by Steinitz’s theorem. For example, all known 1-primitive (*2-simple* in Ziegler’s terminology) 2-simplicial (all 2-faces are triangles) regular partitions of \mathbb{S}^3 are realizable as convex 4-polytopes, which prompted Ziegler (2002) to conjecture that this is always the case.

Remark 15 *Theorem 13 holds if manifold is replaced with homology manifold (see e.g. Seifert, Threlfall for definition and examples).*

Proof. Let $(\mathcal{M}, \mathcal{P})$ be a homology manifold and let the cells of \mathcal{P} be homology balls. Because of the local convexity at the $(n-3)$ -faces, $\mathcal{M} \setminus \text{Sk}_{n-4}(\mathcal{M})$ is actually a manifold. The inductive argument goes without changes, only every time we establish local convexity at a face F of dimension $k \leq n-4$ we also prove that \mathcal{M} is a manifold at all points of F . Thus, \mathcal{M} is a manifold and r is convex. ■

6 Convexity Checker for PL-hypersurfaces

In this section we present a polynomial-time algorithm for checking the convexity of any PL-realization $r : \mathcal{M} \rightarrow \mathbb{R}^n$ ($n \geq 3$) of a regular cell-partition \mathcal{P} of a connected compact $(n-1)$ -manifold \mathcal{M} . The map r under testing is assumed to be dimension-preserving (see Section 2), which implies that each cell $C \in \mathcal{P}$ is homeomorphically mapped by r to an open subset of a subspace of dimension $\dim C$. We do not assume that the realization is an immersion: if it is not an immersion, the algorithm will detect this. We do not make any generic position assumptions.

In describing the algorithm we assume that certain combinatorial and geometric information is readily available. This input information is exactly what should be kept by a convex hull “builder” if it is to use our verification procedure. Later we discuss the complexity of extracting the necessary input information from PL-surface descriptions given in some typical formats.

If any of the subprocedures return false (which corresponds to a detected violation of local convexity), the main procedure returns false as the final answer. The idea of the algorithm is to check that the immersion and the local convexity properties hold at each corner. Recall that the star of a cell C consists of all cells C' with the property that the closure of C' contains C . For each corner C this check is reduced, roughly speaking, to the verification of convexity of a certain cone $K = K(r, \text{Star } C)$ in $r(C)^\perp \cong \mathbb{R}^3$, which is constructed from the poset of $\text{Star } C$ and the restriction of $r : \mathcal{M} \rightarrow \mathbb{R}^n$ to $\text{Star } C$. The cone K completely describes the geometry of the r -realization of $\text{Star } C$ near $r(C)$. Such a cone is not unique – any non-degenerate linear transformation of $K(r, \text{Star } C)$ is just as good as $K(r, \text{Star } C)$. This

reduction from the star of a corner to a cone in \mathbb{R}^3 is done by the subroutine **Reduce-to-3D** called from the procedure **Corner-Checker**.

Let $X \subset \mathbb{R}^n$; then $\text{aff } X$ is its affine span and $\overrightarrow{\text{aff } X}$ the *linear* subspace $\{x-x' \mid x, x' \in \text{aff } X\}$. If V is a set of vectors in \mathbb{R}^n , then we use $\mathbb{R}\langle V \rangle$ or $\text{lin } V$ to denote the linear span of V and $\mathbb{R}_{\geq 0}\langle V \rangle$ the convex cone spanned by V . For a (not necessarily convex) polytope $P \subset \mathbb{R}^n$ with a face F an *affine inner normal* (or, inner skew-normal) to P at F is any vector \mathbf{n} in $\overrightarrow{\text{aff } P}$ such that 1) $\dim(\text{lin } \mathbf{n} \cap \overrightarrow{\text{aff } F}) = 0$ and 2) for any x in the relative interior of P there exists an $\varepsilon > 0$ such that $x + \varepsilon \mathbf{n}$ is in the relative interior of P . We call \mathbf{n} an *inner normal* if, in addition, $\mathbf{n} \perp \text{aff } F$ (i.e. the usual Euclidean normal).

6.1 Input conventions

Mathematically, the input is given as follows:

- (1) the subposet $\mathcal{P}[n-3, n-2, n-1] \subset \mathcal{P}$ of corners, ridges, and facets where it is known in advance which are which;
- (2) an inner normal to R at C for each ridge-corner incidence (R, C) ;
- (3) an inner normal to F at C for each facet-corner incidence (F, C) .

The data in (1) will be referred to as *combinatorial*, and that in (2) and (3), as *linear-algebraic*. We assume that each vector in the linear-algebraic data "knows" the corresponding abstract cells in \mathcal{P} , and that each abstract cell in $\mathcal{P}[n-3, n-2, n-1]$ "knows" all normal vectors related to it. The input data-structure can be implemented as a double-linked adjacency list, with appropriate attribute fields for dimensional and linear-algebraic data. Namely, we can create an adjacency list for the (multi-) graph whose vertex set consists of elements of $\mathcal{P}[n-3, n-2, n-1]$ and whose edge set consists of all pairs (C, C') , such that $C \prec C'$ or $C' \prec C$ in $\mathcal{P}[n-3, n-2, n-1]$. When the input is available in this form we say that the input is given in the *normal form*.

In applications a PL-hypersurface is usually specified by a subposet of the face poset, which includes the vertices or the facets or both; it is usually equipped either with the coordinates of vertices or with the equations (or inequalities) for the facets. If the input is given as the poset $\mathcal{P}[0, n-3, n-2, n-1]$, equipped with the coordinates of the vertices, we say that the input is given in the *vertex form*. If the partition \mathcal{P} is a triangulation, then the linear-algebraic data required for the normal form input can be produced from the vertex form input in time, which is linear in $f_{n-3, n-2}$ and polynomial in the total bit size of the input. Assuming n is fixed, if the face numbers of facets of $(\mathcal{M}, \mathcal{P})$ are bounded by a constant (in $f_{n-3, n-2}$), then the linear-algebraic data for the normal input form can be computed from the vertex form input in $O(f_{n-2, n-3})$ arithmetic operations $(+, -, \times)$.

6.2 Preprocessing

By *preprocessing* in the context of problems of verification of geometric properties we mean any computation that does not depend on the geometric realization (in our case r), but only

on the topology or combinatorics of the object (in our case – the pair $(\mathcal{M}, \mathcal{P})$).

Since \mathcal{M} is a manifold, the facets of $(\mathcal{M}, \mathcal{P})$ making contact at a corner are “glued” to each other in a circular fashion; same can be said about the ridges. Furthermore, both facets and ridges are glued around the corner in the alternating fashion: F_0 - R_{01} - F_1 – \dots , etc. More properly, a topologist would say that the “links” (defined via the 1-skeleton of the *dual partition*, a well-known construction going back to H. Poincare: see Seifert & Threlfall, 1980) of the corners are *circles*. These circles can be thought of as polygons whose vertices correspond to the facets of \mathcal{P} and edges to the ridges of \mathcal{P} . For our convexity checker we need to determine a circular order of ridges around each corner (which is unique up to the choice of direction). The circular structure of the stars of $(n-3)$ -cells implies that for each $(n-3)$ -cell C we have $f_{n-3\ n-2}(\text{Star } C) = f_{n-2}(\text{Star } C) = f_{n-1}(\text{Star } C) = f_{n-3\ n-1}(\text{Star } C)$. The last formula implies that for the whole \mathcal{P} we have $f_{n-3\ n-2} = f_{n-3\ n-1}$. To apply our algorithm for different realizations of the same cell-partition of \mathcal{M} it is reasonable to maintain a circular order of ridges and facets around each corner – this is preprocessing. This circular order is encoded by the corresponding wheel graph $W_m = W_m(C)$ (see Figure 10). The inner normals to facets at $r(C)$ are assigned to the rim edges of W_m and the inner normals to ridges are assigned to the spokes (or rim vertices) of W_m .

The manifold \mathcal{M} may be disconnected, so we have to check for connectivity first. Clearly, \mathcal{M} is disconnected as a topological space if and only if the adjacency graph of the facets, where $(V, E) = (\text{Facets}, \text{Ridges})$, is disconnected. It is well-known that a graph (V, E) can be tested for connectivity, e.g. by the *depth-first* search, in $O(|E|)$ time. Thus, \mathcal{M} can be tested for connectivity in $O(f_{n-2})$ time. From now on we assume that \mathcal{M} has passed the connectivity check.

6.3 Consistency of the input data and guarantees on the output

The algorithm described in the following subsection is guaranteed to work correctly under the following assumptions:

- (1) \mathcal{M} is a connected compact manifold without boundary of dimension $n-1 \geq 2$;
- (2) \mathcal{P} is a finite *regular* cell-partition of \mathcal{M} (i.e. $(\mathcal{M}, \mathcal{P})$ is a regular CW-complex);
- (3) the linear-algebraic data actually comes from some PL-realization r of $(\mathcal{M}, \mathcal{P})$ in \mathbb{R}^n .

6.4 Main Function

Before giving the pseudocode, let us informally describe the working of our algorithm. Recall that to check the convexity of a bounded PL-hypersurface it is sufficient to check the *local convexity* at each corner C , i.e. to check that $r|_{\text{Star } C}$ is a homeomorphism *into* the boundary of a convex body. Based on this idea, our algorithm examines each corner for convexity; these examinations are independent of each other and can be done in parallel. Roughly speaking, locally the star of each corner is supposed to look like the direct affine product of a cone in \mathbb{R}^3 and an $(n-3)$ -subspace. In other words, if we cut the star $r(\text{Star } C)$ by a complementary subspace A of dimension 3, then, in general, we get something that looks like the corner of a

2D PL-surface in $A \cong \mathbb{R}^3$. Alternatively, instead of cutting the star of a corner by an affine 3-subspace, we can get the “3D picture” by projecting the star of the corner along $r(C)$ onto a complementary linear 3-subspace. This “mental visualization” should be taken with some caution, for *geometrically* the corner star may look like a convex surface, but only because the surface has been folded along the ridges so that some facets overlapped – the origami effect.

The main function `Convexity-Checker`, given as Algorithm 1, examines corners for local convexity by calling procedure `Corner-Checker` for each corner. If each corner passes the check by `Corner-Checker`, then `Convexity-Checker` declares the surface $r : \mathcal{M} \rightarrow \mathbb{R}^n$ convex by returning true. On the other hand, if `Corner-Checker` finds out that a corner C is “bad”, it returns the negative verdict false together with a *certificate of failure*. Recall that C can be bad because r is not a local homeomorphism on the star of C , or because there is no convex body such that $r(\text{Star } C)$ lies on its boundary; indeed, a corner can be unlucky enough to fail on both counts. Each of these failure are *violations of convexity*. Thus, when C fails the test, `Corner-Checker` returns (false, not 1-to-1 on $\text{Star } C$) or (false, no convex body for $\text{Star } C$). Note that r may fail to be injective on $\text{Star } C$ because it already failed to be injective on the star of R , where R is one of the ridges of C , – then the certificate has the form “not 1-to-1 on $\text{Star } R$ ”. *The whole program terminates once the very first (in time) violation is detected.*

`Corner-Checker` first checks that r is an immersion on the star of each ridge at C , which is done inside of the “dimension-reducer” `Reduce-to-3D`. If r is an immersion on $\text{Star } C$, then `Reduce-to-3D` reduces the convexity check of the corner to that of its section (see the above paragraph).

In our pseudocode the failure-of-convexity certification is not comprehensive, neither locally nor globally: it tells you just one thing that went wrong. However, even a limited certification is helpful for tasks such as maintaining the convex hull of moving points, or tracing the shape of a simplicial surface whose geometry changes over time (see Guibas, 2004, pp. 1129-1130). However, if complete information about the local convexity and immersion failures is needed, the pseudocode can be easily modified to produce such, at no extra cost (up to a constant factor).

`Convexity-Checker` works on a stack *Corners*, in which we put all corners of $(\mathcal{M}, \mathcal{P})$ prior to starting. The stack is used to simplify the appearance of the pseudocode – any other basic data structure could be used instead. Figure 8 shows the dependency diagram for the modules of our pseudocode: $A \rightarrow B$ means that A may need to call B .

6.5 Corner-Checker

`Corner-Checker` (Algorithm 6.5) tries to reduce the convexity testing at a corner C to that of its projection on a complementary linear subspace, which is spanned by three independent (possibly skew) normals at C to some three facets from $\text{Star } C$ (of course, only if such a triple exists). While in general the projection of $r(\text{Star } C)$ onto a complementary subspace

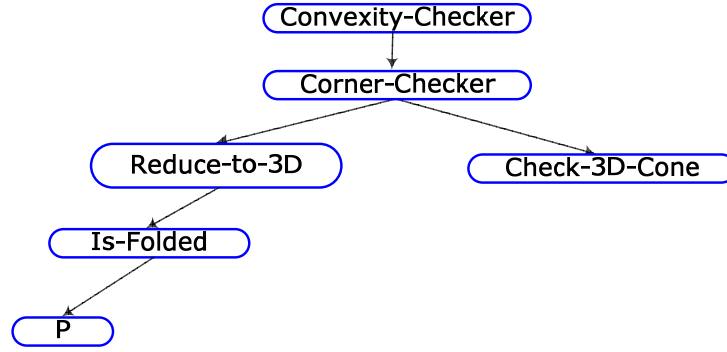


Fig. 8. The dependency graph of the modules

Algorithm 1. Convexity-Checker

Input: poset $\mathcal{P}[n-3, n-2, n-1]$ of corners, ridges, and facets; inner normals to all ridges at all of their corners; inner normals to all facets at all of their corners.

Output: $(answer, certificate)$ where *certificate* specifies the location and reason for failure; if *answer* = true, then *certificate* = none.

```

define global cell[ ] ▷ array of cells; indices start from -1
while stack Corners is non-empty do
   $C \leftarrow Pop(Corners)$ ;  $cell[-1] \leftarrow C$ 
  Encode  $r : Star C \rightarrow \mathbb{R}^n$  as  $(W_m, \mathbf{N})$  ▷ see Subsection 6.6 for details
  for all  $i : 0 \leq i \leq m-1$  do
     $cell[i] \leftarrow$  ridge encoded by vertex  $i$  of  $W_m$ 
  end for
   $(answer, certificate) \leftarrow$  Corner-Checker( $W_m, \mathbf{N}$ )
  if  $answer = false$  then return  $(false, certificate)$  ▷ not convex
  end if
end while
  return  $(true, none)$  ▷ yes, convex

```

looks like a pointed conic surface in \mathbb{R}^3 (e.g. as in Fig. 9, left), it may also look like the star of an edge in a polyhedral surface in \mathbb{R}^3 (Fig. 9, center), or even like a flat 2D-piece in \mathbb{R}^3 , when the surface is flat at the vicinity of $r(C)$ (Fig. 9, right). **Corner-Checker** carefully tests for such degeneracies. Furthermore, the analysis is complicated by the possibility of various local self-intersections; the presence of such self-intersections, when not tested for explicitly, may lead to wrong conclusions.

Corner-Checker uses two major subroutines, **Reduce-to-3D** (Algorithm 3) and **Check-3D-Cone** (Algorithm 6). Recall that the star of any corner C has an intrinsic circular structure, which, for $n = 3$, has the geometric meaning of the order in which flat 2D pieces are glued to each other to form the surface of a polyhedral cone in \mathbb{R}^3 . **Reduce-to-3D** is trying to find three independent vectors among ridge-corner normals at $r(C)$ such that $\overrightarrow{\text{aff } r(C)}$, together with




Fig. 9. The star of corner on the left is simple ($m = 3$) and therefore convex. Note that any immersed PL-hypersurface is *always* locally-convex at all points of its ridges (center) and facets (right).

these vectors, span \mathbb{R}^n . Once three such vectors are found the problem is reduced to dimension 3 via projection onto the span of these three vectors. **Reduce-to-3D** proceeds by going around C , testing for self-intersections and degeneracies, and returning *output_of_reducer*, which is true when **Reduce-to-3D** manages to prove the locally-convexity at C and false when **Reduce-to-3D** finds a violation of the immersion property at the star of one of the ridges of **Star** C , and therefore at **Star** C . If no self-intersections at the ridge stars are detected, and yet the convexity is not proven, **Reduce-to-3D** returns a 3D reduction of the input data, which serves as the input to **Check-3D-Cone**; the reduced data is a rectilinear realization in \mathbb{R}^3 of the wheel graph (see below) that encodes the combinatorics of **Star** C . The (limited) convexity-verification ability of **Reduce-to-3D** is just a byproduct of checking for self-intersections at the level of ridge stars. Namely, **Reduce-to-3D** confirms the local convexity at C only if, in addition to the immersion property, (i) the star of C has at most three ridges, or the geometry of $r(\text{Star})$ is degenerate, i.e. when (ii) $r(\text{Star } C)$ is an $(n - 1)$ -flat or (iii) when it looks like two $(n - 1)$ -flats joined together at an $(n - 2)$ -flat – in the degenerate cases **Star** C may have more than three $(n - 1)$ -cell. The *geometry* of each of the three case is shown in Figure 9 (the flat pieces around the corner may consist of the images of many cells of \mathcal{P}).

Before we come to the details, we need to define a couple of auxiliary notions. For $m \geq 2$ the graph with vertex set $V = \{-1, 0, \dots, m - 1\}$ and edge set $E = \{(-10), \dots, (-1 m - 1)\} \cup \{(01), \dots, (m - 1 0)\}$ is called the m -wheel graph and denoted by W_m (see Fig. 10). Vertex -1 is called the center of W_m ; vertices $V_{rim}(W_m) = \{0, \dots, m - 1\}$ are called rim vertices and edges $E_{rim}(W_m) = \{(01), \dots, (m - 1 0)\}$ rim ridges; edges $\{(-10), \dots, (-1 m - 1)\}$ are called spokes. It is convenient to identify the rim vertices with elements of $\mathbb{Z}_m = \mathbb{Z}/m\mathbb{Z}$; then going from vertex i to vertex j is encoded by adding $j - i$ to $i \pmod m$.




Fig. 10. Graph W_8 .

Once a corner C is popped from the stack, a pair (W_m, \mathbf{N}) is created. Here, m is the number of ridges meeting at C ($m = f_{n-3 n-2}(\text{Star } C)$). This pair consists of (1) the wheel graph $W_m = W_m(C)$, which describes the combinatorics of $\text{Star } C$ and (2) an array of vectors $\mathbf{N}(C)$, whose elements are the inner normals to the r -images of ridges and facets of $\text{Star } C$ at $r(C)$. The center of W_m encodes C itself, the rim vertices encode the ridges of $\text{Star } C$, the rim edges encode the facets of $\text{Star } C$, the spokes encode the corner-ridge incidences; $[0, \dots, m-1]$ is the circular order of the ridges at C determined by the topology of $(\mathcal{M}, \mathcal{P})$.

If **Corner-Checker** returns “false”, it also provides a certificate of violation. The certificate is of the form “not 1-to-1 on $\text{Star } C$ ”, or “not 1-to-1 on $\text{Star } R$ ” (where R is a ridge of $\text{Star } C$), or “no convex witness for $\text{Star } C$ ”. None of these reasons excludes the others.

6.6 Reduce-to-3D

The input to **Reduce-to-3D** is the star $\text{Star } C$ of a corner C together with its realization $r|_{\text{Star } C}$. Formally, the input is a pair $(W_m; \mathbf{N})$, where W_m is the wheel graph encoding the circular structure of $\text{Star } C$, and \mathbf{N} is the array of inner normals to the realizations (r -images) of ridges $(\mathbf{n}_0, \dots, \mathbf{n}_{m-1})$ and facets $(\mathbf{n}_{01}, \dots, \mathbf{n}_{m-10})$ at $r(C)$. \mathbf{N} can also be thought of as a map $\mathbf{N} : E_{\text{rim}}(W_m) \cup V_{\text{rim}}(W_m) \rightarrow \mathbb{R}^n$. Note that we cannot drop facet normals from the input, e.g., in Figure 11 the left and right realizations have the same wheel graph and identical inner ridge normals at C .

Figure 11 (left) shows the star of a vertex that does not violate convexity and immersion assumptions, although the red cell is not convex; on the other hand, the realization on the right has self-intersections, although it does lie on the boundary of a convex body and all cells are convex (two triangles and a pentagon, in each case).

Algorithm 2. Corner-Checker

Input: W_m : wheel graph ($m \geq 2$); \mathbf{N} : non-zero vectors in \mathbb{R}^n indexed by the rim vertices and edges of W_m .

Output: $(bool, cert)$, where $bool$ is “true” or “false”, and where $cert$ is the certificate of violation when $bool = \text{“false”}$ and “none” otherwise.

```

output_of_reducer  $\leftarrow$  Reduce-to-3D( $W_m, \mathbf{N}$ )
if output_of_reducer = true then return (true, none)
else
  if output_of_reducer = (false,  $i$ ) then return (false, not 1-to-1 on Star cell[ $i$ ])
  else
    if Check-3D-Cone(output_of_reducer)=(true, none) then return (true, none)
    end if
  end if
end if

```




Fig. 11. Two flat realizations of a simple star ($n = 3$).

Reduce-to-3D is trying to find three vectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ from the circular list $\{\mathbf{n}_0, \dots, \mathbf{n}_{m-1}\}$ of corner-ridge normals, which span a 3-subspace complementary to $\overline{\text{aff } r(C)}$. If successful, it projects all other corner-ridge normals along $r(C)$ onto this subspace. During the search for three “good” corner-ridge normals Reduce-to-3D is also testing for self-intersections. Since we assume the closures of all cells are realized homeomorphically as polytopes of corresponding dimensions, self-intersections on the level of $r|_{\text{Star } R}$, where R is a ridge at C , can only occur when the images of two facets meeting at R overlap. We call an overlap between two adjacent (the circular order around C) facets *folding* (Fig. 11 shows an example of folding). Some of testing for such overlaps is done inside of Reduce-to-3D and its subroutine Is-Folded, but only until a good triple $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ of corner-ridge normals is found or an overlap is detected, while the rest of testing for self-intersections is delegated to Check-3D-Cone. To detect folding we need to use the *corner-facet normals*, as explained in the beginning of Section 6.6. Once three good vectors are found, all other corner-ridge normals are projected on their span. Self-intersection not detected by Reduce-to-3D are detected by Check-3D-Cone. Note that it is possible, in general, that the angle between two adjacent corner-ridge normals \mathbf{n}_j and \mathbf{n}_{j+1} is greater than π (this angle is determined by \mathbf{n}_{j+1}) and yet the surface is convex at $\text{Star } C$ (as in Fig. 11 left). However, this can only happen when the surface is flat at $\text{Star } C$. Reduce-to-3D guarantees that the output fetched to Check-3D-Cone has the following property: the angle between every two adjacent (in the circular order) vectors $p(j)$ and $p(j+1)$ is less than π (here $p(j)$ is a vector in $\mathbb{R}\langle \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3 \rangle$ representing $j \in V_{\text{rim}}(W_m)$).

Let $p : V(W_m) \rightarrow \mathbb{R}^3$ be a mapping of the vertex set of W_m into \mathbb{R}^3 ; then p is called a *realization* of W_m in \mathbb{R}^3 . If $m \geq 3$, then we can extend W_m to a simplicial complex $[W_m]$ by “filling in” all 3-cycles $(-1 \ i \ i+1)$ (where $i \in \mathbb{Z}_m$). Let us identify $[W_m]$ with the geometric complex in \mathbb{R}^2 where the center of $[W_m]$ is at the origin and the rim of $[W_m]$ is the regular convex polygon, as in Figure 10. We assign to each 2-simplex $(-1 \ i \ i+1)$ a triangle in \mathbb{R}^3 with the vertices $p(-1)$, $p(i)$, and $p(i+1)$. Therefore, the map $p : V(W_m) \rightarrow \mathbb{R}^3$ extends to a map from $[W_m]$ to \mathbb{R}^3 and produces a simplicial surface (with boundary). With a slight abuse of terminology we will say that $p : V(W_m) \rightarrow \mathbb{R}^3$ (where $m \geq 3$) is convex if the extension of p to the 2-complex $[W_m] \subset \mathbb{R}^2$ is a convex surface with boundary (we use p for the extended map as well). While we may encounter W_2 , we will not have a need to associate surfaces in \mathbb{R}^3 with them, for this case is dealt with completely by Reduce-to-3D.

Reduce-to-3D returns (false, i), if r is not an immersion on the star of ridge $\text{cell}[i]$, and “true”, if it is able to verify that r is a convex immersion on $\text{Star } C$. The last scenario is possible

when *geometrically* $r(\text{Star } C)$ near $r(C)$ looks like a corner of a n -simplex, or as a ridge point of a n -simplex, or as a facet of an n -simplex (in which case it is flat). See Figure 9 for 3D illustrations, which can be interpreted as sections as well. In all other cases **Reduce-to-3D** passes the reduced information to the 3D convexity checker **Check-3D-Cone**, which assumes that its input is a 3-dimensional cone, possibly with self-intersections. **Reduce-to-3D** ensures that the cone has at least four rays of which three are linearly independent.

6.6.1 Is-Folded

Let $\mathbf{v}, \mathbf{u}, \mathbf{w}$ be three coplanar *non-zero* vectors. The ordered triple $(\mathbf{v}, \mathbf{u}, \mathbf{w})$ defines a plane angle at the origin in the following way: \mathbf{v} and \mathbf{w} span the two extreme rays of the angle, while \mathbf{u} is an interior vector of the angle – i.e., the function of \mathbf{u} is to specify which of the two open subsets defined by \mathbf{v} and \mathbf{w} is interior to the angle. We denote such angle by $\langle \mathbf{v}|\mathbf{u}|\mathbf{w} \rangle$. Note that $\langle \mathbf{v}|\mathbf{u}|\mathbf{w} \rangle = \langle \mathbf{w}|\mathbf{u}|\mathbf{v} \rangle$.

Is-Folded takes as input a 5-tuple of coplanar *non-zero* vectors $(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e})$, where $\mathbf{a}, \mathbf{b}, \mathbf{c}$ are pairwise distinct and $\mathbf{c}, \mathbf{d}, \mathbf{e}$ are pairwise distinct (these assumptions are made to simplify the pseudocode – the input to **Is-Folded** is guaranteed to satisfy them). **Is-Folded** returns true if the interiors of angles $\langle \mathbf{a}|\mathbf{b}|\mathbf{c} \rangle$ and $\langle \mathbf{c}|\mathbf{d}|\mathbf{e} \rangle$ overlap and false otherwise. For example, Figure 12 shows the case where “folding” takes place: angle $\langle \mathbf{c}|\mathbf{d}|\mathbf{e} \rangle$ “folds over” the angle $\langle \mathbf{a}|\mathbf{b}|\mathbf{c} \rangle$.




Fig. 12. Angles $\langle \mathbf{a}|\mathbf{b}|\mathbf{c} \rangle$ and $\langle \mathbf{c}|\mathbf{d}|\mathbf{e} \rangle$ overlap.

In the pseudocode of this procedure we will use a boolean predicate $P(\mathbf{v}|\mathbf{u}_1, \mathbf{u}_2|\mathbf{w})$, which is defined for any 4-tuple of coplanar vectors $\mathbf{v}, \mathbf{u}_1, \mathbf{u}_2, \mathbf{w}$, where \mathbf{v} and \mathbf{w} are distinct and $\mathbf{u}_i \neq \mathbf{v}, \mathbf{u}_i \neq \mathbf{w}$ for $i = 1, 2$. $P(\mathbf{v}|\mathbf{u}_1, \mathbf{u}_2|\mathbf{w})$ is false if $\langle \mathbf{v}|\mathbf{u}_1|\mathbf{w} \rangle \neq \langle \mathbf{v}|\mathbf{u}_2|\mathbf{w} \rangle$ (Figure 13, left) and true otherwise (Figure 13, right). Intuitively, P returns true if \mathbf{u}_1 and \mathbf{u}_2 both define the same angle with respect to the rays spanned by \mathbf{v} and \mathbf{w} .

The following algorithm shows how to compute $P(\mathbf{v}|\mathbf{u}_1, \mathbf{u}_2|\mathbf{w})$ via standard linear algebra. Recall that an orientation of the real plane \mathbb{R}^2 is an equivalence class of ordered bases, where two bases are equivalent if and only if the change of basis matrix has positive determinant. For any ordered pair of vectors $[\mathbf{a}, \mathbf{b}]$, such that $\{\mathbf{a}, \mathbf{b}\} \subset \text{span}\{\mathbf{v}, \mathbf{u}_1, \mathbf{u}_2, \mathbf{w}\}$, we use $\text{sgn}[\mathbf{a}, \mathbf{b}]$ to denote the orientation of $[\mathbf{a}, \mathbf{b}]$ with respect to some fixed orientation of $\text{span}\{\mathbf{v}, \mathbf{u}_1, \mathbf{u}_2, \mathbf{w}\}$.




Fig. 13. Left: $P(\mathbf{v}|\mathbf{u}_1, \mathbf{u}_2|\mathbf{w}) = \text{false}$. Right: $P(\mathbf{v}|\mathbf{u}_1, \mathbf{u}_2|\mathbf{w}) = \text{true}$.

Algorithm 3. Is-Folded

Input: $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e} \in \mathbb{R}^n \setminus \mathbf{0}$, where $\text{rank}\{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e}\} = 2$, $|\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}| = 3$, $|\{\mathbf{c}, \mathbf{d}, \mathbf{e}\}| = 3$

Output: boolean

```

if  $\mathbf{a}$  and  $\mathbf{e}$  define the same ray then
  if  $P(\mathbf{c}|\mathbf{b}, \mathbf{d}|\mathbf{a}) = \text{true}$  then return true
  else return false
  end if
end if
if  $P(\mathbf{a}|\mathbf{b}, \mathbf{e}|\mathbf{c}) = \text{false}$  and  $P(\mathbf{c}|\mathbf{d}, \mathbf{e}|\mathbf{a}) = \text{true}$  and  $P(\mathbf{b}|\mathbf{c}, \mathbf{d}|\mathbf{e}) = \text{true}$  then return false
else return true
end if

```

Algorithm 4. $P(\mathbf{v}|\mathbf{u}_1, \mathbf{u}_2|\mathbf{w})$

Input: $\mathbf{v}, \mathbf{u}_1, \mathbf{u}_2, \mathbf{w}$ non-zero coplanar vectors with $\text{rank}\{\mathbf{v}, \mathbf{w}\} = 2$

Output: boolean

```

 $s \leftarrow \text{sgn}[\mathbf{v}, \mathbf{w}]$ 
if  $\text{sgn}[\mathbf{u}_1, \mathbf{w}] = \text{sgn}[\mathbf{v}, \mathbf{u}_1] = s$  then
  if  $\text{sgn}[\mathbf{u}_2, \mathbf{w}] = \text{sgn}[\mathbf{v}, \mathbf{u}_2] = s$  then return true
  else return false
  end if
else
  if  $\text{sgn}[\mathbf{u}_2, \mathbf{w}] = \text{sgn}[\mathbf{v}, \mathbf{u}_2] = s$  then return false
  else return true
  end if
end if

```

Algorithm 5. Reduce-to-3D

Input: W_m : m -wheel graph; $\mathbf{N} : E_{rim}(W_m) \cup V_{rim}(W_m) \rightarrow \mathbb{R}^n \setminus \mathbf{0}$

Output: one of $\{\text{true}; (\text{false}, \text{cert}); (W_m, p)\}$

Here $\text{cert} \in V_{rim}(W_m)$, $p : V(W_m) \rightarrow \mathbb{R}\langle \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3 \rangle$ is a “3D picture”, the result of reduction

```

1:  $\mathbf{e}_1 \leftarrow \mathbf{n}_0, \mathbf{B} \leftarrow \{\mathbf{e}_1\}, \mathbf{e}_2 \leftarrow \mathbf{n}_1$            ▷  $\mathbf{B}$  is a maximal ind. set of vectors in  $r(C)^\perp$ 
2: if  $m = 2$  then                                           ▷ there are only two ridges at  $C$ 
3:   if  $(\text{rank}\{\mathbf{n}_0, \mathbf{n}_1, \mathbf{n}_{10}\} = 2)$  and  $(\text{sgn}[\mathbf{n}_0, \mathbf{n}_{01}] = \text{sgn}[\mathbf{n}_0, \mathbf{n}_{10}])$  then return (false, 1)
4:   else return true
5:   end if
6: end if
7:  $i \leftarrow 1$                                            ▷ rim vertex counter mod  $m$ 
8: while  $i \neq 0 \pmod m$  and  $|\mathbf{B}| < 3$  do                 ▷ testing of overlaps and growing  $\mathbf{B}$ 
9:    $i \leftarrow i + 1$ 
10:  if  $\text{rank}\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{n}_i\} \leq 2$  then                 ▷  $\mathbf{n}_i, \mathbf{e}_1, \mathbf{e}_2$  are all in one plane
11:    if  $\mathbf{n}_i = \lambda \mathbf{e}_1$  or  $\mathbf{n}_i = \lambda \mathbf{e}_2$  for some  $\lambda > 0$  then return (false,  $i$ )
12:    end if
13:    if  $|\mathbf{B}| = 1$  then           ▷  $\mathbf{e}_1 = \mathbf{n}_0$  and  $\mathbf{e}_2 = \mathbf{n}_1$  are collinear and contra-oriented
14:      if  $\text{sgn}[\mathbf{n}_0, \mathbf{n}_{01}] = \text{sgn}[\mathbf{n}_0, \mathbf{n}_{0m-1}]$  then return (false, 0)
15:      else
16:        if  $m = 3$  then return true
17:        else  $\mathbf{e}_2 \leftarrow \mathbf{n}_2, \mathbf{B} \leftarrow \mathbf{B} \cup \mathbf{e}_2$            ▷ use  $\mathbf{n}_2$  for  $\mathbf{e}_2$ 
18:        end if
19:      end if
20:    else                                           ▷ in this case we know  $|\mathbf{B}| = 2$ 
21:      if  $\text{Is-Folded}(\mathbf{e}_1, \mathbf{n}_{01}, \mathbf{e}_2, \mathbf{n}_{i-2i-1}, \mathbf{n}_{i-1}) = \text{true}$  then return (false,  $i$ )
22:      end if
23:    end if
24:  else  $\mathbf{e}_3 \leftarrow \mathbf{n}_i, \mathbf{B} \leftarrow \mathbf{B} \cup \{\mathbf{e}_3\}$            ▷ finally we have 3 ind. normals to ridges at  $C$ 
25:  end if
26: end while
27: if  $|\mathbf{B}| < 3$  or  $m = 3$  then return true           ▷  $r(\text{Star } C)$  looks like one of Fig. 9 cases
28: else           ▷ unless  $r(\text{Star } C)$  is flat, angles  $\geq \pi$  between  $\mathbf{n}_j$  and  $\mathbf{n}_{j+1}$  imply non-convexity
29:   for  $j = 0$  to  $m - 1$  do
30:     if  $\mathbf{n}_{j+1} \notin \mathbb{R}_{>0}\langle \mathbf{n}_j, \mathbf{n}_{j+1} \rangle$  then
31:       return (false,  $i$ )
32:     end if
33:   end for
34: end if
35:  $p(-1) \leftarrow (0, 0, 0)$            ▷ center of the wheel is put into the origin
36: for  $j = 0$  to  $m - 1$  do
37:    $p(j) \leftarrow (\mathbf{e}_1 \cdot \mathbf{n}_j, \mathbf{e}_2 \cdot \mathbf{n}_j, \mathbf{e}_3 \cdot \mathbf{n}_j)$ 
38: end for
39: return  $(W_m, p : V(W_m) \rightarrow \mathbb{R}\langle \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3 \rangle)$ 

```

Implementation remark: If Reduce-to-3D is to be implemented for simplicial hypersurfaces for inputs given in the vertex form, then one should not compute Euclidean corner-ridge normals from vertex coordinates when $n \geq 4$. In this case procedures **Convexity-Checker** and **Reduce-to-3D** should be modified so that instead of looking for 3 Euclidean corner-ridge normals that span a 3-subspace complimentary to $\overline{\text{aff } r(C)}$, we find 3 *affine* inner normals of the form $\mathbf{v}_R - \mathbf{v}_0, \mathbf{v}_{R'} - \mathbf{v}_0, \mathbf{v}_{R''} - \mathbf{v}_0$, where $\mathbf{v}_0 \in \mathbb{R}^n$ is a (geometric) vertex of $r(C)$ and $\mathbf{v}_R, \mathbf{v}_{R'}, \mathbf{v}_{R''}$ are geometric vertices of three ridges at C , with the same property. Then the sign (orientation) tests in **Reduce-to-3D** that involve 1,2, or 3 vectors will involve $n - 4, n - 3$, and $n - 2$ vectors respectively, i.e. these tests will become *relative* to some $(n - 3)$ -frame spanned by vertices of $r(C)$.

6.7 Procedure Check-3D-Cone

Informally speaking, we test the surface $([W_m], p)$ for convexity by going around the wheel and checking whether three sign conditions are satisfied or not (certain non-degeneracy conditions should hold as well).

Recall that the sign (orientation) of a list $[\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$ of three vectors in \mathbb{R}^3 is the sign of the determinant of the 3×3 matrix whose i -th row is \mathbf{v}_i . The sign is denoted by $\text{sgn}[\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$. If $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^n$ with $n > 3$, then to define the sign of the triple $[\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$ we need to fix an orientation in the linear span $\text{lin}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$, express the \mathbf{v}_i 's in terms of a basis of $\text{lin}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$, and compute the determinant of the resulting matrix.

Running Frames Conditions:

Set $s := \text{sgn}[p(0), p(1), p(2)]$ and assume $s \neq 0$.

- (1) $\text{sgn}[p(0), p(1), p(i)]$ must be s or 0 for every $i : 2 \leq i \leq m - 1$;
- (2) $\text{sgn}[p(0), p(j), p(j + 1)]$ must be s or 0 for every $j : 1 \leq j \leq m - 2$;
- (3) $\text{sgn}[p(k), p(k + 1), p(k + 2)]$ must be s or 0 for every $k : 1 \leq k \leq m - 1$.

The first condition ensures that all vectors lie on the same side of the plane spanned by $p(0)$ and $p(1)$. Together with the other conditions it guarantees not only that the image $p([W_m])$ is lying on the boundary of a convex body, but also that the surface $p : [W_m] \rightarrow \mathbb{R}^3$ is embedded.

The correctness of **Check-3D-Cone** hinges on the following lemma.

Lemma 16 *Let $p : V(W_m) \rightarrow \mathbb{R}^3$ be a realization of the m -wheel graph ($m \geq 3$) which maps -1 (the center) into $\mathbf{0} \in \mathbb{R}^3$. Suppose the induced map $p : [W_m] \rightarrow \mathbb{R}^3$ on the simplicial 2-complex $[W_m]$ is a homeomorphism on the star of each edge of $[W_m]$. Then $p : [W_m] \rightarrow \mathbb{R}^3$ is an embedding onto the boundary of a convex cone if and only if Conditions (1)–(3) hold.*

Proof. Since $p : [W_m] \rightarrow \mathbb{R}^3$ is a homeomorphism on the star of each edge of $[W_m]$, there exist three vertices of W_m that are mapped to three independent vectors. For the sake of notational convenience we will assume, without loss of generality, that when these vertices are listed in the increasing (circular) order of indices, the sign of the resulting determinant

is positive. We will use the shorthand notation for determinants, i.e. we write $[i, j, k]$ for $\text{sgn}[p(i), p(j), p(k)]$.

If f is a realization of a wheel graph in \mathbb{R}^3 , then denote the predicate “condition (1) holds for p ” by $\text{I}(f)$ and let $\text{II}(f)$ and $\text{III}(f)$ stand for the same predicate, but where condition (1) is replaced with condition (2) and (3) respectively.

If the realization is convex, then Conditions (1)–(3) obviously hold. The proof in the other direction is by induction. For $m = 3$ the result is clearly correct. Suppose the lemma holds for $m = N$. Consider the wheel graph W'_N obtained from W_N by removing the vertex $N - 1$ and connecting the vertices $N - 2$ and 0 by a new rim edge. Denote the resulting map by p' . $\text{I}(p)$ implies $\text{I}(p')$. $\text{II}(p)$ implies $\text{II}(p')$. $\text{III}(p')$ is true if $[N - 1, 0, 1] \geq 0$ and $[N - 2, N - 1, 0] \geq 0$. Observe that $[N - 1, 0, 1] = [0, 1, N - 1] \geq 0$ by $\text{I}(p)$ and $[N - 2, N - 1, 0] = [0, N - 2, N - 1] \geq 0$ by $\text{I}(p)$. Thus $\text{III}(p')$ is true and $p' : [W'_N] \rightarrow \mathbb{R}^3$ is a convex embedding. Therefore, $p : [W_x] \rightarrow \mathbb{R}^3$ is a convex embedding if and only if (A) $p(N - 1)$ lies on the same side of $\text{lin}\{p(0), p(1)\}$ as all other vectors, (B) $p(N - 1)$ lies on the same side of $\text{lin}\{p(N - 2), p(N - 3)\}$ as all other vectors, and (C) $p(N - 1)$ lies on the opposite side of $\text{lin}\{p(0), p(N - 2)\}$ relative to the other vectors.

By $\text{I}(p)$, (A) holds.

By $\text{III}(p)$ we have $[N - 1, N - 3, N - 2] = [N - 3, N - 2, N - 1] \geq 0$. Also, $[0, N - 2, N - 3] \geq 0$ by $\text{II}(p)$. Since p' is convex, we have $[i, N - 2, N - 3] \geq 0$ for all $0 \leq i \leq N - 4$. Thus (B) holds.

By $\text{II}(p)$ we have $[0, N - 2, N - 1] \geq 0$. Also, $[0, 1, N - 2] \geq 0$ by $\text{I}(p)$, so $[0, N - 2, 1] \leq 0$. Since p' is convex, $[0, N - 2, i] \leq 0$ for $1 \leq i \leq N - 3$. Thus (C) holds. ■

The input to **Check-3D-Cone** is a wheel graph W_m ($m \geq 3$), each of whose vertices v is equipped with a corresponding point $p(v)$ in \mathbb{R}^3 . $V(W_m) = \{-1, 0, \dots, m - 1\}$, where -1 is the center of W_m ($p(-1) = \mathbf{0}$), and $[0, \dots, m - 1]$ is a circular order on the corresponding ridges.

Theorem 17 *Under the assumptions stated in Section 6 (Input conventions) **Convexity-Checker** (Algorithm 1) is a correct convexity checker.*

Proof. By Theorem 13 **Convexity-Checker** is correct when the local convexity checks at the corners are correct. The map r is convex at a corner C if and only if

- (1) $r|_{\text{Star } C}$ is an immersion (and, due to our restriction to dimension-preserving PL-realizations, an embedding), and
- (2) its projection on a complementary 3D subspace L is a convex cone.

More rigorously, (2) means that the map $\pi \circ r : \mathcal{S} \rightarrow L$, where \mathcal{S} is a 2-submanifold of $\text{Star } C$ mapped to an affine 3D section of $r(\text{Star } C)$ (as in the proof of Theorem 13) and π is a projection map from \mathbb{R}^n onto L , is an embedding into the boundary of a convex body. Thus,

Algorithm 6. Check-3D-Cone

Input: W_m : wheel graph, $p : V_{rim}(W_m) \rightarrow \mathbb{R}^3 \setminus \mathbf{0}$, $p(-1) = \mathbf{0}$

Output: $(bool, certificate)$.

Below the certificate is given in a simplified form. By examining which of the Running Frames sign conditions failed one can identify, in constant extra time, whether it happened due to the absence of a convex witness, loss of bijectivity, or both.

```

for  $j$  from 0 to  $m - 1$  do                                 $\triangleright j \in \mathbb{Z}_m$ ; all indices are mod  $m$ 
                                 $\triangleright$  checking if the stars of “spokes” are embedded:
  if  $[p(j - 1), p(j), p(j + 1)] = 0$  then
    return (false, not 1-to-1 on Star cell[ $j$ ])
    if  $p(j) \notin \mathbb{R}_{>0}\langle p(j - 1), p(j + 1) \rangle$  then
      end if
    end if
  end for
 $e_0 \leftarrow p(0)$ ,  $e_1 \leftarrow p(1)$ 
 $j \leftarrow 2$ 
while  $j \neq 0 \pmod m$  do                                 $\triangleright$  Running Frames Test
  if  $[e_0, e_1, p(j)] \geq 0$  &  $[e_0, p(j), p(j + 1)] \geq 0$  &  $[p(j), p(j + 1), p(j + 2)] \geq 0$  then
     $j \leftarrow j + 1$ 
  else return (fail, not convex at  $C$ )
  end if
end while

```

our claim is valid if Corner-Checker verifies these two conditions correctly.

Correctness of Check-3D-Cone: We know that p (extended to the 2D simplicial complex $[W_m]$), does not map any two consecutive spokes into the same ray, since $r : \mathcal{M} \rightarrow \mathbb{R}^n$ is a homeomorphism on the closure of each cell C . Check-3D-Cone first checks that the stars of all edges of $[W_m]$ are embedded (on the level of $(\mathcal{M}, \mathcal{P})$ this means that the stars of ridges at C are embedded). Thus, after this check is passed, Lemma 16 is applicable. By this lemma Check-3D-Cone correctly checks that $p : [W_m] \rightarrow \mathbb{R}^3$ is an embedding, which means it correctly checks that $r|_{\text{Star } C}$ is an embedding too. By the same lemma Check-3D-Cone correctly verifies that $p([W_m])$ lies on the boundary of a convex body in \mathbb{R}^3 and, therefore, that $r(C)$ lies on the boundary of a convex body in \mathbb{R}^n . Thus, Convexity-Check is a valid checker. ■

7 Convex hull computation and verification: complexity and robustness

Suppose all computations are done with floating point arithmetic (with `floats`). Although convex hull builders and checkers implemented with `floats` do not guarantee correct outputs, in many situations we can count on certain robustness, which for builders means the output

is close to a convex surface; for checkers robustness means a false positive is only possible when the input surface $(\mathcal{M}, \mathcal{P}; r)$ is in some sense close to a convex surface $(\mathcal{M}, \mathcal{P}'; r')$. We intentionally leave these notions somewhat vague, but it is natural to require that \mathcal{P}' is combinatorially isomorphic to \mathcal{P} almost everywhere and that $r(\mathcal{M})$ is close to $r'(\mathcal{M})$ in \mathbb{R}^n with respect to Hausdorff distance. Furthermore, some uniformity should be required, e.g. we can require that on the subcomplex \mathcal{P}_1 of \mathcal{P} where \mathcal{P} coincides with \mathcal{P}' for each cell $C \in \mathcal{P}_1$ the maps $r|_C$ and $r'|_C$ are uniformly close. A similar notion of robustness makes sense for convex hull builders. In other words, when exact geometry and combinatorics is not that important (as in the case of visualization), convex hull computation and verification with floating point arithmetic is quite meaningful.

More formally, consider the random access machine (RAM) with the unit cost model of computation, where all four arithmetic operations are included into the instruction set. As usual, $f_{i,j}$ denotes the number of incidences between i - and j -faces and $f_i = f_{i,i}$ stands for the number of i -faces. For $n = 3$ if (CO) from Section 1 holds, the algorithms by Mehlhorn et al. (1996) and Devillers et al. (1998), have the same time-complexity of $O(f_0)$. Our algorithm also has the complexity of $O(f_0)$, without requiring (CO) as a precondition; furthermore, its work does not depend on the topology of the input surface. In a more general situation, where conditions (CO) and (S), or one of them, cannot be assumed, the complexity of our algorithm, as well as the earlier algorithms heavily depends on the following factors: i) the combinatorics of the cell-partition (e.g. simplicial or not), ii) the geometry of the realization (e.g. generic positions of the vertices vs. completely general case), iii) the form of the input. Regarding iii), for example, the combinatorial information about the input can be given in the form of the complete poset of faces, or some subposet of faces, such as the vertex-facet graph. Furthermore, certain additional topological information (e.g. the knowledge that the hypersurface is orientable, or a circular order of the facets at each $(n - 3)$ -face) might speed up the convexity verification. The geometry of the realization can be given by the equations of the facets, or by "coherent" inequalities for the facets (CO), or by positions of the vertices, or in the form of inner normals at $(n - 3)$ -cells to the $(n - 2)$ - and $(n - 3)$ -cells.

Now let us consider the general problem of convexity verification. At one end is the simplified setup, where the input hypersurface is simplicial and the realization is sufficiently generic so that the floating point arithmetic can safely be used. Under these assumptions everything is fast, regardless of the method used. At the other end is the completely general setup, where nothing can be assumed. One can also consider "intermediate" models, such as where the input hypersurface is simplicial, but the positions of the vertices are not necessarily generic. Another reasonable assumption would be that the hypersurface is not necessarily simplicial, but the realizations of the facets are known to be convex. We give our algorithm for the most general case, where nothing is assumed. One of the motivations for this generality are the illuminating works of Joswig and Ziegler (2004), Joswig (2004), and Avis, Bremner, and Seidel (1997) who clearly demonstrated that when we cannot assume that the vertices (or hyperplanes) are in general position – or that the dimension of the space is a small fixed number, the convex hull problem is still wide open and interesting from both practical and theoretical viewpoints. As stated by Joswig (2004), "Essentially for each known algorithm there is a family of polytopes for which the given algorithm is superior to any other, and

there is a second family for which the same algorithm is inferior to any other....Moreover, there are families of algorithms for which none of the existing algorithms performs well.” If sufficient linear-algebraic and face incidence information are given about the stars of $(n - 3)$ -faces (see Section 6 for details), the complexity of our algorithmic approach is polynomial in the Turing machine model .

7.1 Complexity analysis

Note that all linear algebra in the algorithm is essentially reduced to comparisons of signs of lists of at most three n -vectors; we will refer to any such calculation as a *sign computation*. Unless mentioned otherwise, as e.g. in the next Subsection, we assume that the input is in the standard form.

- (1) Building the wheel graph for **Star** C takes time linear in the number of ridges of **Star** C .
- (2) Since *Corners* is accessed at most f_{n-3} times, **Corner-Checker** is called at most f_{n-3} times.
- (3) **Corner-Checker**(W_m, \mathbf{N}) requires at most $O(m)$ sign computations.
- (4) **Check-3D-Cone** requires at most $O(m)$ sign computations.
- (5) **Is-Folded** requires a constant number of sign computations.

1 *Suppose the algorithm uses the field algebra $(+, -, \times, \div)$ and each operation has unit cost.* This model is realistic when real computations are conducted with **floats**. If n is fixed, the complexity is $O(f_{n-3} n_{n-2}) = O(f_{n-3} n_{n-1})$. To estimate the complexity in the case where n is one of the parameters describing the input size, we need to estimate the contributions of *sign computations* in (3)–(5). Notice that any sign computation in (3)–(5) deals with, at most, six n -vectors. Since standard linear-algebraic procedures over a field can be used, the complexity of the algorithm is $O(n f_{n-3} n_{n-2})$. For simplicial surfaces this translates into $O(n^3 f_{n-1})$, which is the same as the complexity of Mehlhorn et al.’s algorithm. Their algorithm is faster than ours for some families of non-simplicial PL-hypersurfaces with a very large proportion of simple vertices (where exactly $n + 1$ facets are coming together) under the floating point computing model. Unfortunately, it is next to impossible to enforce vertex-facet incidences while maintaining flatness of non-simplicial facets in computations that use only **floats**. If the geometry of the input surface is sufficiently generic in the sense that at each corner no three ridges lie in the same hyperplane, **Reduce-to-3D** is not needed: any three affine normals can be used for the 3-dimensional reduction. In this case the algorithm requires no divisions, but only evaluation of polynomials of degree at most 3 in the vertex coordinates.

3 *What follows is a discussion of the complexity in the cases where no floating point error can be tolerated.* Let R be the base ring of the computational model: i.e., all numerical input data (such as the coordinates of vertices, the coefficients of normals to $(n - 1)$ -faces etc.) come from R . Furthermore, we assume that $\mathbb{Z} \subset R \subset \mathbb{R}$. When we discuss the degrees of the polynomial predicates evaluated by the algorithm, we consider them as polynomials with integer coefficients in the input parameters. In this context the phrase *arithmetic operation* stands for any ring-theoretic operation $(+, -, \times)$.

Case 1: *the dimension n is fixed.* In this case all linear-algebraic computations can be

done via determinants. Using determinants has an advantage of keeping the degrees of evaluated polynomial predicates as low as possible. Moreover, since in our algorithm the largest determinants are 3×3 , the highest degree of evaluated predicates is 3. Thus, the arithmetic complexity of the algorithm is $O(f_{n-3 n-2})$ and the algorithm evaluates at most $O(f_{n-3 n-2})$ polynomial predicates of degree 3.

Case 2: *the dimension n is not fixed.* If n is not too large, the linear-algebraic computations can still be done via direct determinant evaluations. In each computation we are dealing with at most three n -vectors, which means that we may have to evaluate $\binom{n}{3} 3 \times 3$ determinants to find a minor of maximal rank. Thus, the total arithmetic complexity of the algorithm is $O(n^3 f_{n-3 n-2})$. The case where n is large and the linear-algebraic part of the input is given in the traditional format, i.e. via the vertex coordinates, rather than in the standard form, i.e. via the inner normals to ridges and facets at the corners, is considered in the next subsection.

7.2 Exact Computations over \mathbb{Z}

In here we consider the case of exact computation. The dimension n is not fixed and $R = \mathbb{Z}$. We are interested in the bit complexity of our algorithm, e.g., in the multitape Turing machine model. Since each sign computation involves no more than 6 vectors, the bit complexity of each is $O(nM_b(L))$, if Yap's (2002) ramification of the Bachem-Kannan algorithm is used. Here $M_b(x)$ is the bit-complexity of multiplication of two integers of binary sizes not exceeding x and L is a bound on the binary size of the components of the vectors (see (Yap, 2002) for details). Then the total complexity of the algorithm is $O(nf_{n-3 n-2}M_b(L))$. Devillers et al. (1998) have shown that any convexity checker, *whose work does not depend on the nature of R* , has to evaluate at least one polynomial of degree n – however, this lower bound is mandatory only for those checkers that work the same way for any R . We have:

Theorem 18 *Let $r : \mathcal{M} \rightarrow \mathbb{R}^n$ be a dimension-preserving PL-realization of a manifold $(\mathcal{M}, \mathcal{P})$ of dimension $n - 1$. Suppose the input is in the standard form and all normals have integer coordinates of binary size not exceeding L . There exists a polynomial time algorithm for checking convexity of $r : \mathcal{M} \rightarrow \mathbb{R}^n$ with (multitape Turing machine) complexity of $O(nf_{n-3 n-2}M_b(L))$.*

Now, let us consider the situation where the input is given in the traditional form, i.e. as the poset $\mathcal{P}[0, n - 3, n - 2, n - 1]$, equipped with the coordinates of the vertices. If we have no restrictions on the combinatorics and geometry of the realizations of cells of $(\mathcal{P}, \mathcal{M})$, then it is very difficult, or even impossible, to construct inner normals to facets and ridges at corners from given data. Let us assume that the partition \mathcal{P} is simplicial. In order to do sign computations, we need first write down inner normals for all corner-facet and corner-ridge incidences. For each such incidence we have to deal with roughly n vectors of length n . Computing a Euclidean normal is then reduced to a finding a non-zero solution for a homogeneous system $M\mathbf{x} = 0$ where M is at most n by n matrix. We can use Yap's version of the Bachem-Kannan algorithm to compute the (upper triangular) Hermit Normal Form for the system $M\mathbf{x} = 0$. A non-zero solution vector of at most polynomial size can be found in polynomial time by using standard techniques of linear algebra: we just work our way from

the bottom of the normalized matrix up until all x_i 's are found. Alternatively, one can use a polynomial algorithm in Yap (2002: Sec. 10.8-10.9), based on the repeated application of the Bachem-Kannan algorithm, to further reduce the system to Smith Normal Form and then find a solution. Furthermore, to reduce the complexity, we can deal with each corner C in the following way. If $\mathbf{v}_0, \dots, \mathbf{v}_{n-3}$ are the vertices of $r(C)$, then we first find Hermit's normal form for the matrix $[\mathbf{v}_1 - \mathbf{v}_0, \dots, \mathbf{v}_{n-3} - \mathbf{v}_0]$ and then for each P , where P is a ridge or a facet incident to C , compute an integral normal vector to $r(P)$ at $r(C)$. Then the complexity of all linear-algebraic computations for $\text{Star } C$ is dominated by the complexity of finding Hermit's normal form for the matrix $[\mathbf{v}_1 - \mathbf{v}_0, \dots, \mathbf{v}_{n-3} - \mathbf{v}_0]$, which is $O(n^3 \mathbf{M}_b(L))$ (Yap, 2002). Thus, the total complexity is $O(n^3 \mathbf{f}_{n-3} \mathbf{f}_{n-2} \mathbf{M}_b(L))$ and we have:

Theorem 19 *Let $r : \mathcal{M} \rightarrow \mathbb{R}^n$ be a dimension-preserving PL-realization of a simplicial manifold $(\mathcal{M}, \mathcal{P})$ of dimension $n - 1$. The input consists of the poset $\mathcal{P}[0, n - 3, n - 2, n - 1]$ equipped with the coordinates of the vertices; for each vertex $\mathbf{v} = r(v)$ of $r(\mathcal{M})$ we have $\mathbf{v} \in \mathbb{Z}^n$ and $\|\mathbf{v}\|_\infty \leq 2^L$. Then there exists a polynomial time algorithm for checking convexity of $r : \mathcal{M} \rightarrow \mathbb{R}^n$ with (multitape Turing machine) complexity of $O(n^3 \mathbf{f}_{n-3} \mathbf{f}_{n-2} \mathbf{M}_b(L))$.*

The input requirements in the above theorem can be relaxed. If we know only $\mathcal{P}[0, n - 3, n - 2]$ (or $\mathcal{P}[0, n - 3, n - 1]$) together with the circular orders of facets (or ridges) at all ridges, then $\mathcal{P}[0, n - 3, n - 2, n - 1]$ can be computed at no extra cost.

If $n = 3$, then corners are vertices and the required normals are easy to produce. Now, suppose $n > 3$. How large can the coefficients of the integral normals, discussed above, be? It is obviously possible to produce each such normal as a vector whose coordinates are polynomials of degree at most $n - 3$ in the coordinates of the vertices. Furthermore, Siegel (see e.g. Yap, 2002) proved that a homogeneous system of k linear equations with n variables over \mathbb{Z} has a non-zero solution where each component is bounded in absolute value by $1 + (nA)^{\frac{k}{n-k}}$ (for us $k = n - 3$) where A is the largest of the absolute values of the coefficients. Siegel also showed this bound could not be improved. When n is small enough (say, $n \leq 8$), a vector satisfying Siegel's bound can be found by classical methods of lattice reduction (no efficient methods for finding such a vector are known for large n). If normals satisfying Siegel's bound are used in the algorithmic procedures discussed above, then the largest integers that may appear in sign computations via determinants are of the order $(1 + (nA)^{\frac{n-3}{3}})^3 \leq (1 + \varepsilon)(nA_0)^{n-3}$ - where A_0 is twice the largest of the absolute values of the vertex coordinates and ε is a small positive number.

In practice, a convexity checker based on exact arithmetic can be implemented using one of CGAL's kernels that implements homogeneous coordinates over integers.

7.3 Surfaces in \mathbb{R}^3

The algorithm runs in linear time in the number of vertices when \mathcal{M} is spherical. If f_1 does not exceed $3f_0 - 6$ (the maximal number of edges that a planar graph on f_0 vertices can have), then we can check for connectivity of the input surface in $O(f_0)$ time. However, a sequence of non-spherical PL-manifolds can have the edge number growing quadratically in f_0 . Thus,

it is desirable to check the topological type of the input by just counting 1-cells (edges) in $\mathcal{P}[0, 1]$: once their number exceeds $3f_0 - 6$, we stop and declare the input non-convex. This check, based on that a planar graph on f_0 vertices cannot have more than $3f_0 - 6$ edges, helps preserve the $O(f_0)$ running time bound for PL-surfaces in \mathbb{R}^3 . One may wonder if such a check is necessary, as it seems very likely our algorithm will quickly encounter a non-convex vertex, if the input surface is homeomorphic to a sphere-with-handles or sphere-with-Möbuis-strips. Surprisingly, Betke & Gritzmann (1984), proved that any orientable non-spherical connected 2-manifold can be PL-embedded into \mathbb{R}^3 so that it has exactly 5 non-convex vertices but no fewer! The problem of determining the minimal possible number of non-convex vertices in a PL-immersion of a non-orientable compact 2-manifold is open.

For $n = 3$ the requirements on the combinatorial part of the input can be somewhat relaxed: in what follows we show it is sufficient to know only $\mathcal{P}[0, 1]$, which is the 1-skeleton graph of $(\mathcal{M}, \mathcal{P})$. First, the planarity of this graph can be checked in $O(f_0)$ time. For a planar graph we can also determine the faces (in the combinatorial sense) in linear time – i.e., in $O(f_0)$ time we can create the face-nodes, where each face-node is double-linked to its edge-nodes (e.g. Mehlhorn and Näher, 2000, p. 507). Once we know the faces in terms of their edges, we can double link each face-node to the vertex-nodes of all of its edges. Because of the sphericity of $(\mathcal{M}, \mathcal{P})$ the latter task takes $O(f_0)$ time. Thus, the adjacency list representing $\mathcal{P}[0, 1, 2]$ can be constructed from the adjacency list representing $\mathcal{P}[0, 1]$ in $O(f_0)$ time.

The case of \mathbb{R}^3 is a rather special one. First, $n = 3$ is the smallest dimension for which the techniques of this paper apply. Second, even in the case of \mathbb{R}^n testing convexity for each corner is reduced to testing convexity of a section of the star of this corner, which is essentially equivalent to testing convexity of a cone in \mathbb{R}^3 . In applications a PL-surface in \mathbb{R}^3 is normally specified by its combinatorics and the coordinates of the vertices or equations for the facets: it is therefore important to specify how our algorithm can be applied when the input is given in the traditional form. Namely, suppose we are given $\mathcal{P}[0, 1, 2]$ equipped with the coordinates of the vertices $\mathbf{v}_1 = r(v_1), \dots, \mathbf{v}_{f_0} = r(v_{f_0})$. The corner-ridge normals are then just vectors $\mathbf{v}_i - \mathbf{v}_j$. The question remains how to find corner-facet normals, i.e. vectors pointing from the vertices of the facets into the interiors of the facets. This is easy if it is known that the facets are convex. Otherwise we have the following algorithmic problem. Let C_k be the k -cycle graph. Consider a rectilinear embedding r of C_n into a plane $A \subset \mathbb{R}^3$ – the pair (C_k, r) defines a 2-polytope $P(C_k, r) \subset A$ whose boundary is $r(C_k)$ (here C_k is regarded as PL-manifold). Let v be a vertex of C_k . The problem is to find a non-zero vector $\mathbf{n} \in \vec{A}$ such that $r(v) + \varepsilon \mathbf{n}$ lies in the interior of $P(C_k, r)$. This problem can be solved in time $O(k)$; solving this problem for all facets will require $O(f_0)$ ring-arithmetic operations. Thus, there is no difference in time-complexity between the standard and traditional forms of the input for $n = 3$. We will now restate the observations made in this section in the following theorem.

Theorem 20 *Let $r : \mathcal{M} \rightarrow \mathbb{R}^3$ be a dimension-preserving PL-realization of a 2-manifold $(\mathcal{M}, \mathcal{P})$. Suppose we are given the 1-skeleton of $(\mathcal{M}, \mathcal{P})$ equipped with the coordinates of the vertices. Let L be the upper bound on the bit sizes of the coordinates of the vertices. There exists an algorithm for checking convexity of $r : \mathcal{M} \rightarrow \mathbb{R}^n$ with (multitape Turing machine)*

complexity of $O(f_0 M_b(L))$.

8 Conclusions

This paper describes a local approach to convexity verification of PL-hypersurfaces. The main theoretical result is a characterization of global convexity of an immersed PL-hypersurface in \mathbb{R}^n in terms of the local convexity properties of the surface at its $(n-3)$ -faces. Building on this approach we give a polynomial-time convexity checking algorithm that can be applied for essentially any hypersurface. A better understanding of convexity made it possible to introduce meaningful local certificates of violation; the terms in which previous verification methods described the failure of convexity are of very global character and are not intrinsically linked to the poset structure of the underlying partition. In our method, without increase in asymptotic complexity, we can tell exactly at what ridges and corners convex witnesses do not exist, or the local homeomorphism condition fails (see Section 6). The approach presented in this paper can be generalized to piecewise-polynomial surfaces of small degree (i.e. 2 or 3), where local convexity testing on each simplex can be done via the use of derivatives.

9 Acknowledgements

This research was partially supported by NSF and DARPA under the grant DMS-0310589. The paper has greatly benefited from two detailed reports of an anonymous reviewer.

References

- D. Avis, D. Bremner, and R. Seidel, How good are convex hull algorithms, *Computational Geometry*, (1997) **7**, 265–301, (1997).
- A. D. Alexandrov: *The intrinsic geometry of convex surfaces*, OGIZ, (Moscow and Leningrad, 1948) [Russian]. Translated as *A.D. Alexandrov Selected Works. Part II. Intrinsic Geometry of Convex Surfaces*, Chapman & Hall / CRC, (2006).
- U. Betke and P. Gritzmann, Polyedrische 2-Mannigfaltigkeiten mit wenigen nicht-konvexen Ecken [Polyhedral 2-manifolds with few nonconvex vertices], *Monatsh. Math.* (1984) **97**, No. 1, 1–21.
- M. Blum and S. Kannan, Programs that check their work, In: *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, (1989), 86–97.
- Yu. D. Burago, S. Z. Shefel, The geometry of surfaces in Euclidean spaces. *Geometry*, III, 1–85, 251–256, *Encyclopaedia Math. Sci.*, 48, Springer, Berlin, 1992
- O. Devillers, G. Liotta, F. Preparata, and R. Tamassia, Checking the convexity of polytopes and the planarity of subdivisions, *Computational Geometry*, (1998) **11**, 187–208.
- L. Guibas, Modeling Motion, in *Handbook of Discrete and Computational Geometry*, Chapman & Hall/CRC, pp. 1117-1134, (2004).

- P. Hachenberger, L. Kettner, K. Mehlhorn, Boolean Operations on 3D Selective Nef Complexes: Data Structure, Algorithms, Optimized Implementation and Experiments, *Computational Geometry* (2007), doi: 10.1016/j.comgeo.2006.11.009
- L. Jonker, R. Norman, On locally-convex Manifolds, *Canad. J. Math.* (1973) **25**, 531–538.
- M. Joswig, Software, in *Handbook of Discrete and Computational Geometry*, Chapman & Hall/CRC, pp. 1425–1433, (2004).
- M. Joswig, G. Ziegler, Convex Hulls, Oracles, and Homology, *J. Symbolic Computation* (special issue for ICMS 2002) (2004), **38**, 1247–1259.
- J. Van Heijenoort, On locally-convex Manifolds, *Communications on Pure and Applied Mathematics*, (1952) vol. **V**, 223–242.
- L. Kettner, Using generic programming for designing a data structure for polyhedral surfaces. *Comput. Geom. Theory Appl.*, (1999) **13**:65-90.
- A.M. Kostrikin, Yu.I. Manin: *Linear Algebra and Geometry*, Gordon and Breach, (1989)
- A. Kuzminykh, *Convex bodies with paradoxical topological properties in Hyperbolic space*. 2005 Summer Conference on Topology and its Applications July 10-14, (2005), Denison University, Granville, Ohio.
- K. Mehlhorn and S. Näher: LEDA: A Platform for Geometric and Combinatorial Computing, Cambridge Univ. Press. (2000)
- K. Mehlhorn, S. Näher, M. Seel, R. Seidel, T. Schilz, S. Schirra, U. Stefan, and C. Uhrig, Checking Geometric Properties or Verification of Geometric Structures, In: *Proc. of 12th Annual ACM Sympos. Comput. Geom., Computational Geom.* (1996), 159–165.
- K. Mehlhorn, S. Näher, M. Seel, R. Seidel, T. Schilz, S. Schirra, U. Stefan, and C. Uhrig, Checking Geometric Properties or Verification of Geometric Structures, *Computational Geometry*, (1999) **12**, No. 1-2, 85–103.
- R. T. Rockafellar: *Convex Analysis*, Princeton Univ. Press, (Princeton, New Jersey, 1997).
- K. A. Rybnikov and T. Zaslavsky, Cycle Criteria for Balance in Abelian Gain Graphs with Application to Piecewise-Linear Geometry, *Discrete Comp. Geometry* (2005), **34**, No. 2, 251–268.
- H. Seifert and W. Threlfall: *A Textbook of Topology*, Academic Press (New York, N.Y., 1980) (translated from *Lehrbuch der Topologie*, Leipzig, 1934, by M. A. Goldman).
- G. Ziegler, Face numbers of 4-polytopes and 3-spheres. In: *Proceedings of the International Congress of Mathematicians, Vol. III, Beijing, 2002*, 625–634, Higher Ed. Press, (Beijing, 2002).
- C. Yap: *Fundamental Problems in Algorithmic Algebra*, Oxford Univ. Press, (2000).