

Regular Meshes from Polygonal Patterns

AMIR VAXMAN, Utrecht University

CHRISTIAN MÜLLER, TU Wien

OFIR WEBER, Bar-Ilan University




Fig. 1. Aesthetic and symmetric meshes computed by our method. *Left:* A designer manipulates a $(8, 4)^2$ pattern with a small number of positional constraints (green), obtaining a result where the pattern is locally symmetric. *Center and right:* Existing designs are optimized for pattern quality, revealing interesting geometries. *Center:* Every face is as regular as possible, and the entire mesh appears as a crease folding. *Right:* A hexagonal mesh is optimized for regularity with a fixed boundary. The result is smooth, and the element shapes are uniform.

We present a framework for designing shapes from diverse combinatorial patterns, where the vertex 1-rings and the faces are as rotationally symmetric as possible, and define such meshes as regular. Our algorithm computes the geometry that brings out the symmetries encoded in the combinatorics. We then allow designers and artists to envision and realize original meshes with great aesthetic qualities. Our method is general and applicable to meshes of arbitrary topology and connectivity, from triangle meshes to general polygonal meshes. The designer controls the result by manipulating and constraining vertex positions. We offer a novel characterization of regularity, using quaternionic ratios of mesh edges, and optimize meshes to be as regular as possible according to this characterization. Finally, we provide a mathematical analysis of these regular meshes, and show how they relate to concepts like the discrete Willmore energy and connectivity shapes.

CCS Concepts: • **Computing methodologies** → **Shape modeling; Mesh models; Mesh geometry models;**

Additional Key Words and Phrases: Polygonal patterns, regular meshes, Möbius transformations, architectural geometry

ACM Reference format:

Amir Vaxman, Christian Müller, and Ofir Weber. 2017. Regular Meshes from Polygonal Patterns. *ACM Trans. Graph.* 36, 4, Article 113 (July 2017), 15 pages. <https://doi.org/10.1145/3072959.3073593>

This work is supported by the Israel Science Foundation, under grants 1869/15 and 2102/15, and by the Austrian Science Fund (FWF): P 29981-N35 and I 706-N26.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

0730-0301/2017/7-ART113 \$15.00
<https://doi.org/10.1145/3072959.3073593>

1 INTRODUCTION

Artists and designers constantly explore new types of meshes with desired aesthetics and function. The properties of these meshes are determined by both the combinatorics (valences of vertices and faces, connectivity, and topology), and the geometry (vertex positions). The geometry is important to faithfully represent an underlying surface, and the combinatorics is important for the quality of the representative mesh. Meshes are often designed for realization; for instance, as architectural structures, artistic constructions, and panelled buildings. As such, they need to have original and daring shapes, while adhering to constraints such as prescribed boundaries, limitations on face shapes, and more.

In this work, we compute an “ideal” geometry for a given combinatorial pattern. It is often the case where the pattern is the center of attention in a designed shape, while the geometry is chosen to highlight the beauty of the pattern. Notable contemporary examples are the Soumaya museum [Museo Soumaya], the Kredod [KREOD], the works by Rinus Roelofs [Rinus Roelofs], the Eye Film Museum in Amsterdam [Eye], and De Blob in Eindhoven [de Blob].

Triangle meshes are ubiquitous in computer graphics, but are not very suitable for realization, because of high node complexity. Quad meshes [Pottmann et al. 2007] and hexagonal meshes [Li et al. 2015] gained recent attention in the field of design. In [Jiang et al. 2015], the repertoire was extended to *semi-regular* tilings. We extend the range even further by incorporating Archimedean, aperiodic, and hyperbolic patterns into design.

The question is then what the ideal geometry is for a given pattern. Common and natural measures of aesthetics are symmetry and smoothness, and by them we define the concept of *regular meshes*, where either the faces, or the vertex neighborhoods of a pattern, are as rotationally symmetric as possible. For this purpose, we define




Fig. 2. We compute a regular $(4, 3, 4, 3, 3)$ mesh for prescribed handles, as the centerpiece of an architectural design.

novel energies that encode the symmetry and anisotropy of 1-rings and faces. By minimizing our energies for prescribed positional constraints, we obtain as-regular-as-possible meshes. Fig. 1 and 2 depict examples of various meshes designed with our method, in the context of architectural geometry [Pottmann et al. 2015].

We compute meshes that minimize these energies, and therefore the mesh elements of our optimal results are “as-regular-as-possible”. The energies are based on manipulating three basic quantities: the vertex based quaternionic cross-ratio [Hertrich-Jeromin 2003], a novel 1-ring based polygon that encodes the regularity of this 1-ring, and a face based normal ratio, with which we control the regularity of faces. In addition, we show how existing definitions for ideal geometry from combinatorics, such as connectivity shapes, are related to our framework. The meshes that our method computes are aesthetic, and enrich the vocabulary of mesh design considerably. Finally, our regularity definitions correlate with notions in discrete differential geometry, like the discrete Willmore energy [Bobenko and Schröder 2005]. Our contributions are:

- We introduce a new method to design and optimize meshes to become as regular as possible, allowing for mesh editing with arbitrary boundary and handle-based positional constraints.
- We introduce two novel measures of regularity: one measures the rotational symmetry of faces (Euclidean regularity), and one measures symmetry of mesh elements invariantly under Möbius transformations (Möbius-regularity).
- We define geometric characterizations of regularity, based on an analysis of quaternionic ratios on the elements of the mesh.
- We generalize and optimize for regularity for all mesh combinatorics, including pattern imperfections, singularities, and unordered or irregular patterns.
- Our definitions relate to the discrete Willmore energy.

2 RELATED WORK

Geometry to/from combinatorics. There are several common mesh-design paradigms. A popular choice is *combinatorics-from-geometry*, where a given shape is reassigned a connectivity with a different resolution, quality (e.g., more regular elements), or different types of elements (e.g., a triangle mesh is converted to a quad mesh) [Botsch et al. 2010, Chap. 6]. Conversely, in the *geometry-from-combinatorics* approach, the vertex positions are inferred from the connectivity [Bobenko et al. 2006; Isenburg et al. 2001]. This approach is desired when a specific pattern is to be realized (for instance, a quadrilateral grid), and the surface is not known in advance; rather, it is designed to adhere to positional and structural constraints. This process is referred to as *form-finding* [Tang et al. 2014]. Our work follows the latter type of approach, where we optimize for a mesh that is the most regular, or symmetric, for a prescribed pattern.

Polygonal mesh design. Most methods are not designed to deal with polygonal meshes of arbitrary connectivity, and are not susceptible to such an adaptation. A possible way to construct polygonal meshes with rotationally-symmetric faces is by using methods that perform general projections [Bouaziz et al. 2012; Tang et al. 2014], but there is no insight for what an optimal solution is, or whether it is possible to obtain one. In addition, they do not define notions of symmetries in 1-rings and vertices. We nevertheless offer a modification to [Bouaziz et al. 2012] in Section 5. In [Jiang et al. 2015], polyhedral patterns were computed for given surfaces, optimizing for planarity, while keeping the faces symmetric. This requires the decomposition of a surface into strips, and thus they target patterns of specific symmetries, with limited generality and control over the behavior of singularities. As such, the quality of the result depends heavily on feature alignment and combinatorics. Instead, in our work, the geometry is created from any underlying polygonal pattern, and does not require any specific structure to work. We do not adhere to a given embedding of a prior shape, but rather compute the geometry that would optimize the quality of the mesh elements.

Critical meshes. Other works created geometry from constraints and combinatorics, with well-known energies, for which the resulting surfaces were critical points. A prominent example is that of discrete minimal surfaces [Pinkall and Polthier 1993]. They have been defined for triangular meshes, in the language of finite-elements (FE), and for quadrilateral meshes, general polyhedral surfaces, and hexagonal meshes [Bobenko et al. 2006, 2010; Müller and Wallner 2010; Pottmann et al. 2007] in the language of discrete differential geometry (DDG). DDG differs from FE, in the sense that it attempts to create discrete definitions that share properties with continuous concepts, rather than discretize continuous quantities in a piecewise-continuous way. The definitions in our paper follow the DDG philosophy, by defining ratios on elements of a mesh. Minimal surfaces are a subclass of constant mean curvature surfaces. Their aesthetics, and their elegant mathematical properties, motivated their computation as triangular meshes [Pan et al. 2012], and as quad and hex meshes in [Pottmann et al. 2007].

Other approaches for creating critical shapes from given combinatorics are the creation of minimal surfaces from circle patterns with given combinatorics [Bobenko et al. 2006]. This mathematical result,




Fig. 3. Examples of tilings with regular faces in the plane. A 3-uniform tiling (top-left), semi-regular (Archimedean) tilings (top-center and -right), and a regular tiling of the hyperbolic plane (bottom-left).

however, does not provide a flexible design method of minimal surfaces in practice, for instance to solve boundary value problems. Two works are closely related in context to our work. In [Bobenko and Schröder 2005], a discrete Willmore energy is devised by looking at the intersection angles of circumscribed circles of the triangles of a mesh. Then, a mesh minimizing the Willmore energy is computed. However, they do not target mesh regularity, but rather only local cosphericality. We show that our characterization of regularity is related to the discrete Willmore energy in Section 6, and argue that this is why our results exhibit a smooth “as-spherical-as-possible” appearance, while optimizing for the regularity of the elements. Isenburg *et al.* [2001] create geometry for a given mesh connectivity by optimizing for edges of unit lengths, which is closely related to our Euclidean regularity energy, as we discuss in Section 7.

3 REGULAR MESHES

We work with meshes $\mathcal{M} = \{\mathcal{V}, \mathcal{E}, \mathcal{F}\}$. A face $f \in \mathcal{F}$ is a closed (not necessarily planar) polygon with an arbitrary number of edges (we refer to this number as the *valence* of the face). We denote meshes and patterns where all faces have the same valence as *pure meshes* (e.g., pure quad meshes, where all faces have valence 4) and meshes with diverse face valences as *mixed meshes*. Consequently, we denote the neighborhood, or 1-ring, of a vertex $v \in \mathcal{V}$ as *pure*, if all the adjacent faces have the same valence (e.g., any vertex in a pure triangle mesh), and we denote it as *mixed* otherwise.

A *pattern* is the combinatorics of a mesh, and it is mostly defined with some repetition, or periodicity. A *tiling* is a realization of a pattern (assigning coordinate positions to vertices) such that the realized mesh covers a subset of some space. Notable examples are planar and spherical tilings (see Fig. 3). The terms “pure” and “mixed”, in the sense of combinatorics, apply here as well. The classical definition of tiling requires the realization to respect the repetition of the pattern. For instance, that faces of the same valence are congruent. Examples range from the simplest pure tilings in the plane, to intricate tilings combining regular and rhombic faces interchangeably, like Penrose tilings [Grünbaum and Shephard 1986]. Repetitive combinatorial patterns are usually identified by enumerating the valences of faces around each vertex, and using upper indices for multiples if all 1-rings are identical up to Euclidean transformations. For instance, a pure-quad pattern is 4^4 .




Fig. 4. Perfect Euclidean-regular meshes. Top row: Icosahedron (left), Archimedean solid (rhombicosidodecahedron; center), PolyCube mesh (right). Bottom row: developable quad mesh (left), two Archimedean solids glued together (center), developable triangle strip (right).

Planar tilings. There are three *regular tilings* of the plane where all faces are the same regular polygon: triangle, quad or hexagon. They are pure regular tilings, according to our valence characterization. There are some other notable tilings with regular faces. They are often denoted as *semi-regular* or *Archimedean*. These patterns are mixed and regular by our definition. The set of tilings with regular faces is however larger than all the above mentioned (semi-)regular patterns. Their generalizations can be further identified, e.g., by their invariance under transformations of the wallpaper group. See Fig. 3 for examples of planar tilings. The guiding principle is that the angle sum around a vertex in a planar tiling must be 2π , or otherwise not all faces can be regular.

Spherical tilings. Regular and semi-regular tilings on the sphere correspond to Platonic and Archimedean solids. Each face is a regular n -gon. The five platonic solids are pure regular tilings of the sphere. Any two vertex stars of an Archimedean solid are congruent, and in the case of Platonic solids the vertex stars are even rotationally symmetric (with rotations of $2\pi/\text{valence}$).

Hyperbolic tilings. There are also well-known regular and semi-regular tilings in the hyperbolic plane. Any two regular n -gons are related by a hyperbolic isometry (i.e., by a Möbius transformation in the Poincaré disk model). The heptagonal tiling, for example, does not exist in the Euclidean plane, as the angle sum of three regular heptagons around a vertex (Fig. 3 bottom-left) equals $\frac{15}{7}\pi > 2\pi$.

Inspired by the aesthetics of (semi-) regular tilings, we aim to create realizations for meshes with a similar appearance. For a flexible design process, we allow a designer to freely prescribe positional constraints. In order to understand the results that we would get from the design process, we need to explore and expand the definition of regularity for general meshes, as we do in the following.

3.1 Euclidean-regular meshes

We define a mesh, pure or mixed, to be perfectly *Euclidean-regular* if all its faces are perfectly regular in the usual sense (i.e., each face is rotationally symmetric). Well-known examples of Euclidean-regular




Fig. 5. Meshes optimized for Euclidean regularity. The quad mesh assumes perfect developability, the heptagonal mesh assumes a hyperbolic pattern, as the prescribed angle defect is negative, and the hexagonal mesh has the appearance of a dual geodesic dome.

meshes are Platonic and Archimedean solids, and all the aforementioned regular and Archimedean planar tilings. Other interesting examples are that of PolyCube meshes [Tarini et al. 2004], and developable meshes with perfect faces. This definition also includes meshes that are composed by gluing these well-known examples together. Some of these meshes are depicted in Fig. 4.

By the assumption that every face is perfectly Euclidean-regular, it is then evident that the combinatorics determines the inner angles of the faces completely, which in turn determine the Gaussian curvature at each vertex (i.e., their angle defect). Since all faces of a perfectly Euclidean-regular mesh are regular polygons, all edges must have the same length. Hence, once an edge length is selected, the geometry of perfectly Euclidean-regular meshes is determined from the combinatorics up to isometric deformation. In fact, in many cases, the only possible isometric deformation is a rigid motion. A notable exception is the case of perfectly Euclidean-regular meshes with 4^4 combinatorics (pure quads): we obtain developable (and flexible) surfaces. In general, more degrees of freedom are possible for meshes with boundaries: for instance, a strip of perfect triangles. See Fig. 4 for such examples.

In general, not every pattern admits a perfectly-regular realization. Moreover, additional user specified positional constraints further restrict the solution space. In light of this, perfectly Euclidean-regular meshes are for the most part restricted and inflexible with respect to design needs. Our method then targets the generation of *as-Euclidean-regular-as-possible* meshes. We denote them as just “Euclidean-regular” meshes in short (omitting “perfectly”). We define them as meshes of arbitrary given combinatorics, where the faces are optimized to be close to perfectly regular faces, under given constraints (such as boundary or positional handles).

Even though such meshes are at best approximately regular, we can predict the geometric realization by noting that perfectly-regular faces prescribe a specific metric and specific Gaussian curvature: For example, modeling in our context (a part of the mesh) with the combinatorics of a *planar* regular pattern will therefore result in a mesh which is (in that part) “as developable as possible”, and a mesh optimized for Euclidean-regularity with a regular pattern from the hyperbolic plane will assume the shape of a negatively curved surface in \mathbb{R}^3 (i.e., hyperbolic surface points). See Fig. 5 for examples.




Fig. 6. Perfect Möbius-regular meshes. The truncated dodecahedron ($10^2, 3$) pattern (left), and a pure triangle patch (right) under Möbius transformations. The meshes, as a result, are embedded in spheres.

Knowing this, a designer who optimizes for Euclidean regularity, can design the combinatorics in advance to fit her expectations.

3.2 Möbius-regular meshes

We offer an additional definition of regularity, extending our repertoire of regular meshes. We define a mesh to be *perfectly Möbius-regular* if every 1-ring can be individually transformed by a Möbius transformation to a perfectly Euclidean-regular respective 1-ring, embedded in a generalized sphere (i.e., sphere or plane). As a consequence, every face is Möbius equivalent to a regular face. We call the original perfectly-Euclidean regular 1-ring the *canonical embedding*. Note that this definition categorically excludes rings that are not embeddable in a sphere with perfectly Euclidean-regular faces. For instance, 1-rings with negative Gaussian curvature. Example: a junction of 5 quads in a PolyCube mesh is not a canonical embedding for Möbius regularity, despite having Euclidean-regular faces.

Broadly speaking, perfectly Möbius-regular meshes are spherical meshes (in a Möbius-geometry sense, including planar). Such meshes either are Archimedean, or stereographic projections of regular planar tilings (Fig. 6). Spherical regular tilings, like Archimedean solids, are then perfectly Euclidean- and Möbius-regular.

Our motivation for introducing the definition of Möbius regularity is that, from a design perspective, they provide more options for results that are beautiful and symmetric. The shapes of the regular elements transform into aesthetic shapes, as a Möbius transformation is *conformal*. One should note that Euclidean regularity is *not* a special case of Möbius regularity. The aforementioned PolyCube junctions are one example, but non-planar developable meshes are another, as the vertices are not Möbius regular. These two notions of regularity are related and intersecting, yet distinct.

Nevertheless, as in the case of perfectly Euclidean-regular meshes, the flexibility gained by the Möbius-equivalence between elements is limited if we only consider perfectly Möbius-regular meshes. Thus, we focus on geometric realizations that are *as-Möbius-regular-as-possible*, for given constraints. Again, we denote these as Möbius-regular meshes (without ‘perfectly’). We can predict their appearance from the combinatorics itself: they are always as-spherical-as-possible (including planes), where additionally the elements are as-regular-as-possible. See Fig. 7 for some examples. The notion of trying to become as-spherical-as-possible brings up the concept of the Willmore energy [Bobenko and Schröder 2005]. Indeed, as we show in Section 6, perfectly Möbius-regular meshes are related to discrete Willmore surfaces.

3.3 Boundary behavior

The definition of Möbius regularity readily extends to boundaries: a boundary vertex is Möbius regular if it has a canonical embedding as a planar boundary vertex, where all faces are regular. This includes spherical patterns, as boundary vertices can always be embedded in the plane without distortion. The definition of Euclidean regularity is thus straightforward.

This is considered as the *natural* boundary behavior, and mimics the effect of cutting out faces from a wallpaper pattern, seamlessly along the edges (see inset). As a result, the boundary curvature is also predetermined (for the canonical embedding), and different patterns exhibit considerably different boundaries accordingly (see Fig. 7).

For instance, it is expected to find smooth boundaries when the angles add up to π , and jagged boundaries otherwise. We employ natural boundary conditions in all examples in our paper, wherever they are not explicitly constrained as handles.






Fig. 7. Examples of Möbius-regular meshes. The Eye mesh is as spherical as possible with the given two boundaries. The three planar patterns, where faces are color-coded according to valence, are optimized with the same positional constraints (black spheres), and natural boundary conditions. The resulting shapes are similar on the inside, and considerably different on the boundary, where the natural curvature of each pattern is obtained.

3.4 Imperfect patterns

Not all mesh combinatorics admit a perfectly-regular realization (Euclidean or Möbius), even without considering any positional or boundary constraints. For Euclidean regularity, every mesh that cannot be optimized to developable patches without cutting, or changing the combinatorics, falls into this category (Fig. 8).


A basic example of imperfection for Möbius regularity is this geodesic dome (see inset), i.e., the sphere with 12 cone singularities of index $\frac{1}{6}$. Consider two neighboring vertices with valences 5 and 6 (green and red, resp.). The perfectly Euclidean-regular 1-ring for the vertex of valence 6 is the respective element in a 3^6 tiling.



The perfectly Euclidean-regular 1-ring of valence 5 is a spherical 1-ring, as the angle defect $\frac{\pi}{3}$ is positive. For both 1-rings to be perfectly Möbius-regular, the mutual flap (two dashed triangles in the inset) must be Möbius equivalent to perfect flaps of both valences 5 and 6. This is impossible, as can be easily verified with the invariance of the cross-ratio (see Section 4). Thus, the singular vertices of this geodesic dome cannot be perfectly Möbius-regular without some distortion in the neighboring vertices.

The challenge is in fact more considerable: not every conceivable 1-ring admits a canonical embedding even independently: for instance, there is no canonical embedding for quad-mesh vertices of valence 5, or triangle-mesh vertices of valence 7, due to the negative angle defect. Nevertheless, we want our method, and definition of regularity, to incorporate imperfect combinatorics, so that our method would be general and useful. This is also essential for meshes of higher genus that admit negative-index singularities. We therefore must define a *canonicalization* of imperfect 1-rings.

As a motivational example to how we canonicalize imperfect 1-rings, consider a pure 1-ring of valence d . We then choose the perfectly symmetric solution: arrange the vertices in a regular d -gon, and put the central vertex at the barycenter (see inset). We then consider 1-rings that are Möbius equivalent to this canonical reference as regular.



Several examples of pure regular meshes from imperfect patterns are given in Fig. 8. While we defined a way to work with imperfect pure patterns, we still do not know how to cater to mixed imperfect patterns at this stage, or how to produce these meshes in practice. To do these, one needs to produce a consistent geometric characterization of regularity and canonical 1-rings, that can be optimized for; we do so in the following section.

4 GEOMETRIC CHARACTERIZATION OF REGULARITY

We optimize for pattern realizations that are as regular as possible, either Möbius or Euclidean, where we cater to perfect as well as imperfect patterns automatically. For that, we proceed as follows: First, we introduce a compact characterization of the geometry of perfectly-regular meshes. To this end, we consider three quaternionic ratios of edges on a mesh: the classical cross-ratio, a normal ratio for two consecutive edges within each face, and the corner tangent on every adjacent three vertices within a face. With them, we define canonicalized embeddings to imperfect patterns.

4.1 Preliminaries

Face flaps, 1-rings and boundary polygons. We consider a mesh $\mathcal{M} = \{\mathcal{V}, \mathcal{E}, \mathcal{F}\}$ as before. To define our ratios we refer to several configurations/elements in a mesh: a single *face-flap* contains two faces f, g with a common edge e (see Fig. 9 top-left). A *1-ring* around a central vertex v consists of all faces adjacent to that vertex. In addition, we use the *vertex star* around a central vertex v which consists of the edges vu_1, \dots, vu_n emanating from v and which also contains their endpoints u_1, \dots, u_n in a cyclic order (Fig. 9 right). Vertex stars and 1-rings are identical in pure triangle neighborhoods (Fig. 9 bottom left). Otherwise, the triangles of a vertex star are not necessarily faces of the mesh (see the three dashed triangles in Fig. 9




Fig. 8. Imperfect meshes optimized for regularity. Optimized for Möbius regularity (MR), the genus-0 meshes (Bunny and Fandisk) become spherical with faces that are as-regular-as-possible. The genus-1 Rockerarm interestingly achieves a cyclide appearance, which is the as-spherical-as-possible shape for genus-1. Optimized for Euclidean regularity (ER), Fandisk and Bunny achieve approximate, but imperfect, developable appearance. Color coding is for the energies defined in Section 5.

right). We predominantly use the vertex star in our definitions. We denote the polygon generated by the sequence of the endpoints u_1, \dots, u_n of a vertex-star around v as its *boundary polygon*, and mark it as B_v .

Quaternions. Quaternions are defined as $q = [r, x, y, z] = [r, v] \in \mathbb{H} \cong \mathbb{R}^4$, where $r = \text{Re}(q)$ is the *real part*, and $v = (x, y, z) = \text{Im}(q)$ the *imaginary part*. The noncommutative *product* of two quaternions $[r_1, v_1]$ and $[r_2, v_2]$ is $[r_1 r_2 - \langle v_1, v_2 \rangle, r_1 v_2 + r_2 v_1 + v_1 \times v_2]$. Quaternions q with $\text{Im}(q) \neq 0$ can be uniquely decomposed as $q = s[\cos \phi, \sin \phi v]$, where $s \geq 0$ is the *modulus* (absolute value), $0 \leq \phi < \pi$ is the (principal) *argument*, and $v \in \mathbb{R}^3$ is the (unit) *vector* of the quaternion. Throughout this paper, we mark unit vectors with bold letters. Given two quaternions a, q , we call the product $aq a^{-1}$ a *rotation*, and $aq \bar{a}$ a *similarity*. We avoid using the term “conjugation” for rotations, in the context of this paper, in order to dispel confusion with the conjugate $\overline{[r, v]} = [r, -v]$. Rotations only alter the vector part, while similarities also change the modulus.

We represent vertices and edge vectors in \mathbb{R}^3 by *imaginary quaternions* $q_i = [0, x_i, y_i, z_i]$. The product of two imaginary quaternions is then $q_1 q_2 = [-\langle \text{Im}(q_1), \text{Im}(q_2) \rangle, \text{Im}(q_1) \times \text{Im}(q_2)]$. Representing this product by $s[\cos \phi, \sin \phi v]$, we get that $s = |q_1| |q_2|$, $\phi = \pi - \alpha$, where α measures the convex angle between the vector parts, and $v = (q_1 \times q_2) / |q_1 \times q_2|$ is the normal vector to both quaternions.




Fig. 9. The setting of our exposition. *Top-left*: A face-flap consisting of two triangles adjacent along a common edge. *Bottom-left*: A pure triangle 1-ring around a vertex v , identical to the vertex star of v . *Right*: A mixed polygonal mesh. The vertex star around vertex i consists of all triangles generated by the edges through i (red dashed). The polygon jkl is its boundary polygon.

A quaternionic *Möbius transformation* is of the form $q \mapsto w = (aq + b)(cq + d)^{-1}$. We only regard *imaginary-preserving* transformations, s.t. w is imaginary if and only if q is. Imaginary transformations map \mathbb{R}^3 to \mathbb{R}^3 (with the point at ∞). Properties of imaginary-preserving transformations were explored in [Vaxman et al. 2015].

4.2 Quaternionic cross-ratio

We abbreviate edge vectors between vertices by $q_{ij} = q_j - q_i$. The *quaternionic cross-ratio* of four points q_i, q_j, q_k, q_l is $\text{cr}_q[i, j, k, l] := q_{ij} q_{jk}^{-1} q_{kl} q_{li}^{-1}$. For completeness, we repeat some of its relevant properties from [Vaxman et al. 2015].

Intersection angle. Consider the circumcircles of vertices ijk and kli , as in Fig. 10 (left). The cross-ratio argument $\phi[i, j, k, l]$ encodes the intersection angle $\psi[i, j, k, l]$ between both oriented circles as $\phi[i, j, k, l] = \pi - \psi[i, j, k, l]$. We thus obtain the well known result that q_i, q_j, q_k, q_l are circular (lying on a common circle) if and only if the cross-ratio is real (zero imaginary part).

Under transformation. The edge vector of two transformed points q_i, q_j under a Möbius transformation is obtained by

$$\overline{(cq_i + d)}^{-1} (q_j - q_i) (cq_j + d)^{-1}. \quad (1)$$

Thus, a Möbius transformation, applied to four points q_i, q_j, q_k, q_l , rotates the old cross-ratio cr_q by $\overline{(cq_i + d)}$ to

$$\text{cr}_w[i, j, k, l] = \overline{(cq_i + d)}^{-1} \text{cr}_q[i, j, k, l] \overline{(cq_i + d)}.$$




Fig. 10. *Left*: Triangle flap with intersection angle ψ between the circumcircles of neighboring triangles. *Right*: Imaginary parts of cross-ratios corresponding to two permutations of the four points.

The modulus $|cr|$ and the argument ϕ are preserved.

4.3 The corner tangent

The *corner tangent* at q_i is defined on three consecutive vertices q_k, q_i, q_j within a face, constituting an oriented corner, as follows:

$$t[k, i, j] := -q_{ki}^{-1} q_{jk} q_{ij}^{-1} = q_{ki}^{-1} + q_{ij}^{-1}.$$

The last equality holds since $p^{-1} + q^{-1} = p^{-1}(q + p)q^{-1}$. Note that the corner tangent is always imaginary. The cross-ratio $cr[i, j, k, l]$ can be expressed in terms of corner tangents:

$$cr[i, j, k, l] = t[k, i, j]^{-1} \cdot t[k, i, l].$$

Other than the simple algebraic relation between the corner tangents t and the cross-ratio cr , the corner tangent has the following interesting geometric meaning (see also Fig. 10 left):

LEMMA 4.1. *Consider the circumcircle of q_i, q_j, q_k , oriented according to this defining triangle. Then, $t[k, i, j]$, placed at q_i , is in oriented tangential contact to the circle.*

We include a proof of Lemma 4.1 in Appendix A. This lemma provides a geometric interpretation to the vector part of the cross-ratio: it is orthogonal to the tangents of two intersecting circles ijk and ikl , and therefore it is in the direction of the radius vector of the sphere which is defined by these circles, or equivalently by the four vertices of the triangle flap.

The tangent polygon. By observing the definition of the corner tangents, it is straightforward to see that at a vertex star with central vertex v and vertices u_1, \dots, u_n , we have:

$$\sum_{i=1}^n t[u_i, v, u_{i+1}] = 0,$$

where the indices are taken modulo n . By that property, if we consider the vectors t as edge vectors of an abstract polygon, then this polygon is closed for every vertex star (when v is not a boundary vertex). We denote it as the *tangent polygon* of v (see Fig. 11), and mark it as T_v . The imaginary parts of the cross-ratios $cr[v, u_{i-1}, u_i, u_{i+1}]$ then serve as the normals to each respective corner of T_v . Note that the vector part of a cross-ratio is always in the direction of the cross product $t[k, i, l] \times t[k, i, j]$. As such, if and only if T_v is convex (and planar), all vector parts are co-directional.




Fig. 11. The normal ratios N_j, N_l are defined on corners j, l and their vector parts are orthogonal to corresponding triangles. The vector part of the cross-ratio cr_i of the flap $ijkl$ is orthogonal to the corresponding edges $t[k, i, j], t[l, i, k]$ of the tangent polygon T_i .




Fig. 12. *Left:* Regular vertex star. The cross-ratio of each flap in the 1-ring is $[\cos \phi, 0, 0, \sin \phi]$. *Right:* The vertex star after a Möbius transformation, and its tangent polygon undergoing a similarity.

Under a Möbius transformation, a corner tangent $t[k, i, j]$ transforms via a similarity (cf. Eq. (1)):

$$t_w[k, i, j] = (cq_i + d) t_q[k, i, j] \overline{(cq_i + d)}. \quad (2)$$

Thus, Möbius transformations of the vertex star result in similarities of T_v (see Fig. 12). This invariance of the tangent polygon makes it the centerpiece in our definition of Möbius regularity, and in the canonicalization of imperfect patterns, as we see in Section 4.5.

Ratios on faces. To be able to discuss Möbius regularity of meshes, we also need to discuss ratios on the faces, and not only on vertex stars. Two triangles are always Möbius equivalent to each other. Therefore, we only discuss faces with valence $d \geq 4$ in this context. To be able to characterize faces in a similar way to that of vertex stars, we consider the cross-ratios of each four consecutive points of a face of valence d , with vertices (u_1, \dots, u_d) . If the face is Möbius equivalent to a Euclidean-regular face, the cross-ratio is predetermined by d alone; see Section 4.5 for the specifics.

The tangent polygon and the cross-ratio are objects of Möbius geometry, and are consequently used by us to define Möbius regularity; in the following we define an object of Euclidean regularity.

4.4 Normal ratios

Having defined the tangent polygon, whose edge vectors are the corner tangents, and the cross-ratio of a flap in a vertex star around a given central vertex v , we proceed to define the normal ratio, that pertains to two consecutive edges q_{ij}, q_{jk} within a face. We define the *normal ratio* of each corner q_i, q_j, q_k as (see Fig. 11):

$$N[i, j, k] = q_{ij} q_{jk}^{-1}.$$

If the vertices are clear from the context, we use N_j for the normal ratio instead of $N[i, j, k]$, for brevity. Note that the imaginary part of the normal ratio is always orthogonal to the plane spanned by the two vectors. Its modulus encodes the ratio of their lengths $|q_{ij}|/|q_{jk}|$, and its argument encodes the *convex* angle $\angle ijk$ between them. Unlike the corner tangent and the cross-ratio, the normal ratio is not invariant under Möbius transformations, as $N_w[i, j, k] = (cq_i + d)^{-1} N_q[i, j, k] (cq_k + d)$. The normal ratio transforms as a rotation when the Möbius transformations have no inversion ($c = 0$). As such, the modulus and the argument of normal ratios are invariant to similarities. We also obtain the elegant identity:

$$cr[i, j, k, l] = N_j \cdot N_l.$$

4.5 Quantifying regularity with ratios

Euclidean Regularity. A face of valence d is perfectly Euclidean regular if and only if all the normal ratios equal

$$[\cos(\chi_d), -\sin(\chi_d)\mathbf{n}_f], \quad \text{where } \chi_d = 2\pi/d, \quad (3)$$

and \mathbf{n}_f is a unit normal to the face. This is in fact stating a straightforward definition of Euclidean regularity: in a regular polygon, every edge is a rotation of its adjacent edge by the angle specified by the valence. We use it in this manner for optimization in Section 5.

Möbius regularity of faces. Under a Möbius transformation, the cross-ratios of triangle flaps of vertex stars undergo rotations. Consequently, the moduli and arguments of the cross-ratios are invariant to the transformation. For a perfectly Möbius-regular d -gon (i.e., a d -gon which is Möbius equivalent to a Euclidean-regular d -gon), we get that all cross-ratios of four subsequent vertices are real (since regular d -gons are circular) and equal

$$\text{cr}[u_i, u_{i+1}, u_{i+2}, u_{i+3}] = [-(1 + 2\cos(2\pi/d))^{-1}, 0, 0, 0].$$

The tangent and the boundary polygon. The tangent polygon T_v of a perfectly Euclidean-regular vertex star with edge length 1 equals its boundary polygon B_v (i.e., $T_v = B_v$), since:

$$t[k, i, j] = q_{ki}^{-1} + q_{ij}^{-1} = -(q_{ki} + q_{ij}) = q_{jk}.$$

Möbius regularity of pure vertex stars. Recall that a vertex star is pure when the adjacent faces all have equal valences. Let us consider a perfectly Möbius-regular vertex star of valence n with central vertex v . Its canonical embedding is a perfectly Euclidean-regular star embedded in a sphere, where all edge lengths are $\|v - u_i\| = 1$, and B_v is a regular n -gon. This leads to the following theorems:

THEOREM 4.2. *The tangent polygon T_v of a pure Möbius-regular vertex star of valence n is a Euclidean regular n -gon.*

COROLLARY 4.3. *In a pure Möbius-regular vertex star of valence n the cross-ratios $\text{cr}[v, u_{i-1}, u_i, u_{i+1}]$ of all flaps around the 1-ring are equal to each other and to*

$$[\cos \phi_n, \sin \phi_n \mathbf{n}_v], \quad \text{where } \phi_n = (n-2)\pi/n,$$

where \mathbf{n}_v is a unit vector.

This is easily verified. See Fig. 12 for illustration. Note that n is the degree of the vertex, rather than that of the faces. For instance, in the case of a cube or the dodecahedron, we have $n = 3$. Furthermore, the tangent polygon undergoes a similarity while the mesh undergoes a Möbius transformation, and therefore retains its Euclidean regularity. \mathbf{n}_v can be considered as “the vertex normal” of the vertex star, as it is codirectional with the radius vector of the circumsphere to this star.

Möbius regularity vertex stars in mixed meshes. The characterization of Möbius regularity of pure vertex stars with quaternionic ratios readily generalizes to mixed meshes. Assume that the valences of the faces around a central vertex v are (d_1, \dots, d_n) , and that they are ordered cyclically around the vertex. We refer again to the canonical embedding for the perfectly Euclidean-regular vertex star embedded in a sphere, where edges emanating from v all have the same length 1 and angles α_i of faces i measure $(d_i - 2)\pi/d_i$.




Fig. 13. The tangent polygon, which identifies with the vertex star. Left: the canonical embedding of a mixed (4, 3, 4, 5) pattern. Right: the tangent polygon with the phases of the cross-ratios that are used to quantify ideal regularity of this vertex star.

The identical tangent and boundary polygons for this star have the following elegant theorem (see also Fig. 13):

THEOREM 4.4. *If the vertex star is co-spherical, and the lengths of edges emanating from v are equal to 1, then the associated boundary polygon B_v , and the identical tangent polygon T_v , are concyclic.*

PROOF. Let us denote the sphere that contains the entire vertex star (boundary polygon and vertex v) by S_1 . The boundary polygon is equidistant from v , and thus is contained in its own sphere S_2 for which v is the center. The two spheres are distinct (since $v \in S_1$ but $v \notin S_2$) which implies that the boundary/tangent polygon lies in the intersection $S_1 \cap S_2$ and is consequently concyclic. \square

COROLLARY 4.5. *The tangent polygon of a perfectly Möbius-regular mixed vertex star is concyclic, equal to the boundary polygon of the vertex star in the canonical embedding.*

See Fig. 13 for an example. The cross-ratios of the flaps in the vertex star behave accordingly: they all have the same vector \mathbf{n}_v , and moduli and arguments derived from the ratios of edges in the tangent polygon. Note that pure vertex stars are an interesting subset: tangent polygons of perfectly Möbius regular vertex stars are concyclic, and tangent polygons of pure vertex stars are additionally regular. As we see in Section 6, the weaker property of merely being planar has an important geometric meaning as well.

Computing the tangent polygon. Our analysis of the tangent polygon and cross-ratio for Möbius regular vertex stars has a direct purpose: given a combinatorial pattern, determine the ideal tangent polygon T_v at each vertex star v , so that an embedding of the vertex star for which T_v is the tangent polygon, is then Möbius regular. We then optimize for this property to obtain Möbius-regular meshes.

As we opt for a tangent polygon that is concyclic, there is a simple way to compute the ideal one:

- For each star triangle ijk from a 1-ring face of valence d , calculate the length of the ideal boundary edge l_{jk} according to its valence. This is the length of a diagonal in a unit d -gon.
- Create an embedding of the boundary polygon B_i in a circle with the given lengths. This is also the perfect tangent polygon T_i .

The process is a bit technical, but straightforward; we give exact details in Appendix B. To get the moduli and the arguments of

the cross-ratio, simply compute them from the embedding. For boundary conditions, the tangent polygon is open. Nevertheless, since we assume natural boundary condition, we can predict what the “missing faces” that were supposedly cut out of the boundary are, and collect their lengths to complete the polygon. As an example, a boundary with a quad and two triangles have natural regular angles $\frac{\pi}{2} + 2 \cdot \frac{\pi}{3}$. The missing “phantom” polygon is predicted to have an angle of $\frac{5\pi}{6}$ to complete to 2π .

It may seem that the analysis of the tangent polygon is unnecessary, as we can just build a canonical embedding and use it. However, the insight about the structure of the tangent polygon is essential for the generalization of Möbius regularity into imperfect patterns as in the following, where no such Euclidean embedding exists.

4.6 Regularity in imperfect patterns

Our method is designed to accommodate every given input pattern, and provide as-regular-as-possible realizations. It is expected that we reproduce perfect embeddings for perfect patterns (when positional constraints do not interrupt), but not every pattern admits a canonical embedding, as we discuss in Section 3.4. The face-based quantities, cross-ratios and normal ratios, are agnostic of the imperfectness of the pattern; we try to make each face as regular as possible independently, according to its valence d .

For imperfect vertex stars, we need a consistent *modus operandi*. The practical meaning of our insight about the structure of cross-ratios and tangent polygons, is that we have an intuitive way to canonicalize imperfect vertex stars:

Definition 4.6. A vertex star is canonicalized to be Möbius regular if its tangent polygon is concyclic, with the lengths l_{jk} , ascribing to the valences of the faces in the vertex star.

This definition reproduces the tangent polygons of pure and mixed perfect vertex stars, and in addition the intuitive generalization we gave for imperfect pure vertex stars in Section 3.4. We then just adopt the circle embedding as our *modus operandi*: we don’t have to know if the vertex stars are perfect or imperfect. As a




Fig. 14. Meshes optimized for Möbius regularity with direct optimization. Top: designing a Möbius regular shape from a patch of the (6, 4, 3, 4) pattern by prescribing position constraints (green dots). Bottom: optimizing existing designs. Left: a (12, 6, 4, 6) pattern. Right: a triangle mesh, giving the appearance of a “blob toy”. The colors refer to the E_{MR} energy per vertex.

consequence, we get a single canonicalization of regularity for all patterns, that we optimize for in the following section.

To summarize, our characterization seamlessly caters to any type of input mesh, while including the exact definition of regularity for meshes that do possess a canonical regular embedding.

5 REGULAR MESH DESIGN

We next define the regularity energies, based on our quaternionic ratios. We hereby refer to both canonical and canonicalized embeddings as just “canonical” from this point on. The designer starts from either a generic embedding of a desired pattern (e.g., a patch in the plane), or some existing realization of a pattern, and prescribes positional constraints to direct the design.

5.1 Regularity energies

Möbius regularity. Consider a mixed 1-ring of valence n around vertex w_i . We use unit-length vectors \mathbf{n}_i per vertex to represent the vector parts of the ideal mutual cross-ratios, and define the following energy, combining face-based and vertex-star Möbius regularity:

$$E_{MR} = \sum_{f \in \mathcal{F}} \sum_{p=1}^d \left| \text{cr}[w_p^f, w_{p+1}^f, w_{p+2}^f, w_{p+3}^f] - \left[\frac{-1}{(1+2 \cos(2\pi/d))}, \mathbf{0} \right] \right|^2 + \sum_{w_i \in \mathcal{V}} \sum_{\text{flap}(ijkl)} \left| \text{cr}[w_i, w_j, w_k, w_l] - l_{ijkl} [\cos \phi_{ijkl}, \sin \phi_{ijkl} \mathbf{n}_i] \right|^2,$$

where d is the valence of f , ϕ_{ijkl} and l_{ijkl} are the argument and modulus, respectively, associated with the cross-ratio of the flap, as computed from the concyclic tangent polygon, and w_i^f denotes a vertex i of the face f . $E_{MR} = 0$ only when all cross-ratios are equal to the ideal one, which means every 1-ring and every face are perfectly Möbius-equivalent to their canonical embeddings. Note that the moduli and the arguments of the common auxiliary cross-ratio are constants and pre-computed. The variables in this energy are the vertex positions w_i that constitute the cross-ratios, and the auxiliary normalized vector parts \mathbf{n}_i . By prescribing positional constraints, simply as a substitution into the coordinates of the vertices, the designer can deform any pattern to her will. We show examples of E_{MR} -minimizing meshes in Fig. 14.

Euclidean regularity. Consider a face of valence d . Similar to the Möbius regularity energy, we utilize a unit-length auxiliary vector variable \mathbf{n}_f per face, and define the energy (cf. Eq.(3)):

$$E_{ER} = \sum_{f \in \mathcal{F}} \sum_{\substack{(k,i),(j) \\ \text{adjacent edges} \in f}} \left| w_{ij} w_{ki}^{-1} - [\cos(\chi_n), -\sin(\chi_n) \mathbf{n}_f] \right|^2,$$

to minimize the differences between the normal ratios of the polygons to a central common auxiliary normal ratio per face. Again, both the vertex positions w_i and the normalized normal vectors \mathbf{n}_f are variables in our system. We show examples of meshes optimized for Euclidean regularity in Fig. 15.

Direct optimization. Given coefficients λ_{MR} and λ_{ER} , a direct optimization method minimizes the following unconstrained energy:

$$E_R = \lambda_{MR} E_{MR} + \lambda_{ER} E_{ER}.$$

To get Euclidean regular meshes, we set $\lambda_{MR} = 0$ and $\lambda_{ER} = 1$. Möbius regular meshes often have a degree of freedom, since the results

are unique up to a Möbius transformation. While sufficient positional constraints remove this ambiguity, in practice it is preferable to settle with a small coefficient in ER. Thus, for all our examples we use $\lambda_{MR} = 1$ and $\lambda_{ER} = 0.01$ for Möbius regularity. For initial solution, we use the original positions as the initial w_i , and solve for the initial \mathbf{n}_i and \mathbf{n}_f by minimizing E_R , where the positions are the initial w_i set as constant. These are local linear systems.

We show examples of meshes designed with direct optimization, in Fig. 14 and 15. It is naturally possible to balance the energies differently to get blended regularity, and we depict this in Fig. 16.

5.2 Local-global based approach

The energy is smooth, but nonlinear, as the cross-ratio is a nonlinear function of the vertices and we use the normalized auxiliary cross-ratio vector. As such, it can be solved by methods that target nonlinear least squares, and we use the Levenberg-Marquadt algorithm [Nocedal and Wright 2006]. Nevertheless, every iteration is costly, and computing Jacobians for quaternionic variables is tiresome. It is then worthy to explore an alternative method of computing with local-global alternating projections, and specifically with shape projections [Bouaziz et al. 2012]. Trying such an approach is appealing, as it involves local iterations of simple problems, and efficient global linear solves with a fixed positive-definite matrix, for which a Cholesky factorization could be computed in a preprocessing step. To this end, we need to define an alternative formulation of our energy, which is not equivalent, but which still theoretically optimizes for Euclidean and Möbius regularity:

$$E'_R = \lambda_{MR} \sum_{i=1}^{|\mathcal{V}|+|\mathcal{F}|} |A_i W_i - P_{MR}(A_i W_i)|^2 + \lambda_{ER} \sum_{i=1}^{|\mathcal{F}|} |A_i W_i - P_{ER}(A_i W_i)|^2, \quad (4)$$

where W_i is a subset of vertex positions participating in a projection, i.e., vertex stars for MR and face vertices for either ER or MR, A_i




Fig. 15. Meshes optimized for Euclidean regularity. Top and bottom left: the optimized geometry for these semi-regular patterns. Note that they appear as if made from conic patches, as the original patterns are planar. Bottom right: the optimized dragon exhibits a shape that is made from patches that are as developable as possible, having the appearance of a “paper toy”.




Fig. 16. Balancing regularities. We use $\lambda_{MR} = 1$, and change λ_{ER} according to the specified values. The result balances between a spherical and smooth look (dominant MR), to one that tends to developability (dominant ER).

are averaging operators, and $P_{ER/MR}$ are projection operators that we describe in the following. The formulation is identical to the one given by [Bouaziz et al. 2012], and we refer the reader for exact details of the averaging and the nature of the local-global iterations.

Projection for Euclidean regularity. The core operation in the local iteration is to take the vertices w_i in iteration k , and produce the closest (in the least-squares sense) ideal vertices $P_{ER}(w_i)$. For Euclidean regularity, it’s simply the closest regular polygon. This was studied in [Bouaziz et al. 2012], and we do not alter the formulation.

Projection for Möbius regularity. For faces, one needs to find the closest polygon $P_{MR}(w_i)$ that is related by a Möbius transformation to a canonical regular polygon of the same valence. However, as both the canonical polygon, and the ideal result are planar, we can reduce this problem to two substeps: first, find the global rotation that brings the current polygon W_i as best as possible into the xy plane where the canonical regular polygon is, and subsequently project the coordinates onto the plane (set $z = 0$) to form complex coordinates η_i . Second, find a planar Möbius transformation between the canonical polygon and η by minimizing the following:

$$\sum_{w_i \in W_i} |a\eta_i + b - z_i(c\eta_i + d)|^2,$$

where $a, b, c, d \in \mathbb{C}$ are the variables of the complex Möbius transformation, and z_i are the vertices of the canonical polygon. It is possible to set $a = 1$ to remove the homogeneous ambiguity. The local projection P_{MR} operator comprises of the inverse to this Möbius transformation, followed by the inverse to the rotation above.

Projection for vertex stars is more challenging. We do not always have a canonical embedding to compute a Möbius transformation from. Even if we do, finding the quaternionic transformation in \mathbb{R}^3 directly is not linear anymore, as quadratic imaginary-preserving conditions on a, b, c, d have to be met [Vaxman et al. 2015].

The first challenge we solve is to create a canonical embedding. This is straightforward: we create the concyclic tangent polygon, and put the central vertex in the middle. The resulting tangent polygon is equal to the boundary polygon. For perfect patterns, this is simply the stereographic projection of the canonical embedding to the plane—a Möbius transformation—and we can use this as a canonical embedding. We denote this embedding with vertices q_i .

The second challenge is to find the best-fitting Möbius transformation between q and the current vertex star w . Recall that the tangent polygon undergoes a similarity. We proceed as follows:

[i] Find the similarity between the tangent polygons of the canonical embedding, and that of the current vertex star W_i . This is computing $G_v = (cq_v + d)$, as per Equation (2).

[ii] Solve for c, d by minimizing

$$\sum_{w \in W_i} |w_{v_i}^{-1} - (cq_i + d)(q_v - q_i)^{-1} \overline{G_v}|^2$$

(Eq. (1) in an inverse form). This is a linear quaternionic system. Recall that w is given as the current vertex positions.

[iii] Compute $P_{MR}(W_i)$ by applying the Möbius transformation.

Analysis. The energies of Eq. (4) are not the same as the energies of Sec. 5.1. The main difference is that the formers are no longer scale invariant. Essentially, reducing edge size lowers the energy, which is not the case for the energies of Sec. 5.1; there, every edge is represented in both q_{ij} and the inverse q_{ij}^{-1} in another term, and each term is scale invariant. We make an empirical comparison of both alternatives in Fig. 17. We observe that the direct optimization requires much less iterations to converge, but has more costly iterations. The lightweight iterations of the local-global approach make up for the excessive amount of iterations.

In the case of Euclidean regularity optimization (the “Tent” mesh, $\lambda_{MR} = 0, \lambda_{ER} = 1$), the results are quite similar. On average, we found the entire runtime of the local-global approach to take about 75 – 80% of the total running time per example, comparing to direct optimization. However, the code for both approaches is not optimized for performance, and the local step could be effectively parallelized. In light of this, we recommend the local-global approach for regularity optimization that is dominantly Euclidean-regular.

In the case of Möbius regularity (the “Teddy” example, $\lambda_{MR} = 1, \lambda_{ER} = 0.01$), the local-global approach achieves good results as well as the direct optimization, but different in appearance. However, in other cases (same parameters), the local-global approach tends to converge to local minima. These local minima have an interestingly common “dumbbell shape” (the “3D bar” and the “Bishop”), which is almost perfect everywhere, but for a small “neck”. Direct optimization manages to avoid this type of minima. We conjecture that these minima are related to edge sizes in the local-global formulation, but cannot provide a theoretical proof.

It is worthwhile to inquire whether the regularity energies of Section 5.1 can be formed in a local-global approach directly. For instance, given $cr[i, j, k, l]$, compute vertex positions w_i , and vice-versa (for Möbius regularity). However, the global problem is still nonlinear and difficult, and our experimentations did not show any advantage over the direct optimization. We chose to use direct optimization in all the examples in our paper for uniformity. Nevertheless, we see the potential of using local-global approaches as a good alternative for regularity, with a different formulation for Möbius regularity, but that needs to be studied further.

5.3 Implementation Details

The direct optimization method is a nonlinear least-squares system, with $|V|$ imaginary quaternionic variables, which are $3|V|$ scalars in total. In addition, we need $|F|$ auxiliary imaginary quaternions for \mathbf{n}_f , and $|V|$ more for \mathbf{n}_i . We used the Google Ceres solver [Agarwal et al.], that uses the Levenberg-Marquadt algorithm. Ceres provides two major advantages: first, automatic differentiation alleviates a

somewhat tedious calculation of Jacobians in analytic form. Second, it allows for a parametrization of \mathbf{n}_i and \mathbf{n}_f as unit quaternions, without the need for a unit-length regularizer on our behalf. We used libhedra [Vaxman 2016] for the representation and manipulation of polyhedral patterns. The entire computation for a direct




Fig. 17. Direct vs. local-global optimization. Top, the Teddy mesh with Möbius-regular optimization. Middle, the Tent mesh, with Euclidean-regular optimization. Bottom two: 3D bar quad mesh, and Bishop triangle mesh, exhibiting the “dumbbell” result. Convergence refers to the ratio of norm of the energy (E_R or E'_R), to the initial respective energy, for fair comparison.

optimization is between 3-10 sec. for meshes of about 500-2000 vertices on a 4GHz i7 iMac with 32GB memory. This is not a real-time rate, but still attractive for a design process.

6 ANALYSIS

6.1 Discrete Willmore energy

The Willmore energy for a smooth surface S is defined by $W(S) = \int_S (H^2 - K) dA$, and is a measure for how much S differs from a sphere. We next discuss a discrete version of this energy.

We established in Section 4 that the tangent polygon of a pure vertex star is a regular polygon if and only if the star is Möbius regular, and that it is concyclic with specified lengths for a mixed vertex star. In addition, we mentioned that the vectors of the cross-ratio of a flap $ijkl$ at i and at k , extended as lines, are coplanar and meet at the center of the sphere circumscribing the flap. Finally, Möbius regular vertex stars (mixed or pure) are by definition cospherical. These geometric insights are special cases of a more general result in discrete differential geometry:

THEOREM 6.1. *A vertex star is cospherical (all vertices lie on one sphere) if and only if the vector part of the cross-ratios $cr[v, u_i, u_{i+1}, u_{i+2}]$ at the central vertex v are all parallel.*

COROLLARY 6.2. *The tangent polygon T_v at v is planar if and only if the vertex star of v is cospherical.*

Thm. 6.1 appears in [Gwynne and Libine 2012] in another form, without the relation to our tangent polygon. It interestingly ties directly with the formulation of [Bobenko and Schröder 2005]: they define a *discrete Willmore energy* by

$$W(v) = 2\pi - \sum \phi_{v, u_i, u_{i+1}, u_{i+2}}, \quad (5)$$

where ϕ is the phase of the cross-ratio at vertex v for every flap. Assuming the inner angles of the tangent polygons are $\psi = \pi - \phi$, it is evident that Equation (5) is similar to Cor. 6.2.

Nevertheless, this similarity is not generally an equivalence. The phase of the cross-ratio is *always* defined to be the convex angle between the tangents; this is also used this way in [Bobenko and Schröder 2005], as they define the phase in the range $[0, \pi)$ as well (using the inverse cosine), and optimize for the flat and convex angle sum. However, the tangent polygon can be planar and nonconvex. In that case, Cor. 6.2 holds: the vertex star is cospherical, the cross-ratio vectors are parallel, but not all in the same direction, and Eq. (5) does not hold. That means that not all spherical vertex stars have $W(v) = 0$. Nevertheless, we should note that cosphericality in the discrete case is considerably less intuitive than in the continuous case: a mesh with vertices on the sphere can be very “zigzagged” and crumpled, and planarity of tangent polygons does not guarantee any smoothness of the mesh. For example: take any conceivable mesh, and normalize the coordinates. The mere planarity of the tangent polygon is thus too weak to serve as a quality measure.

Inscribable meshes. This case is acknowledged in [Bobenko 2005; Bobenko and Schröder 2005]. They consider *inscribable* polyhedra, that can be realized as a convex spherical mesh, and that achieve $W = 0$. Noninscribable polyhedra always have $W > 0$, and any spherical embedding of them has planar nonconvex tangent polygons. See Fig. 18 for examples of inscribable and non-inscribable polyhedra.




Fig. 18. Inscribability of meshes. Left: an inscribable mesh, with convex tangent polygons. Right: a non-inscribable mesh, where its spherical embedding shown has nonconvex tangent polygons. It is cospherical, but with $W \neq 0$ according to Equation (5).




Fig. 19. Left: a noisy realization of the small Rhombicosidodecahedron. Right: our optimization restored the perfect state.

Relation to Möbius regularity. Our optimization for *regular* tangent polygons requires a stronger property than planarity. As such, Möbius-regular meshes in general tend to look “as-spherical-as-possible”. However, the strong property of regularity means that Möbius regularity is not a discrete Willmore energy, and may have different minima in some cases. See Fig. 25 for an example.

Conformal Willmore Flow. In [Crane et al. 2013], a *conformal Willmore flow* is introduced, producing a surface that is Willmore critical, with the property of being conformal to the original surface. It is cogent to create a relevant formulation in our method, as a sort of Möbius regularity, when the original surface is the canonical embedding, rather than a perfect 1-ring. Their method is defined as a discretization of a continuous flow, but it is possible to obtain the cross-ratio based quantities with discrete conformality [Springborn et al. 2008]. We will explore this possibility in the future.

7 RESULTS

In this section, we give an empirical analysis of our method by several scenarios, and relate to other works in the literature.

Robustness. In Fig. 19, we show how our method is able to reproduce a small rhombicosidodecahedron, with perfect Möbius and Euclidean regularity, from a very noisy state. This demonstrates how our optimization is able to handle initial solutions far from the




Fig. 20. Meshes with PolyCube combinatorics have perfect Euclidean-regularity that is achieved in our optimization.




Fig. 21. The cube has both a closed-form minimal solution for Möbius regularity (the symmetric spherical shape), and a perfect solution for Euclidean regularity. Our optimization achieved both.




Fig. 22. Convergence under refinement. The mesh converges to a spherical shape with increasingly better regularity everywhere, but, as predicted, errors on vertices around singularities do not vanish.

theoretical minimum. We demonstrate this capability further for meshes with PolyCube combinatorics in Fig. 20, where the perfect Euclidean regular solution is obtained, even for meshes of a higher genus. Finally, in Fig. 21, we show that our optimization obtains the best possible Möbius and Euclidean regularity for the twisted cube. These examples show that our optimization can reach ground-truth results, even with considerably different initial solutions.

Convergence. We show an interesting experiment in Fig. 22, where a pure-quad mesh is repeatedly subdivided by Catmull-Clark, and optimized for Möbius regularity. While all results are close to a symmetric spherical embedding, as expected, the more refined meshes have better error profiles everywhere, but at the vertices adjacent to singularities. This is the expected behaviour for imperfect patterns that we discussed in Sec. 3.4.

Connectivity shapes. Euclidean regularity is strongly related to the concept of connectivity shapes [Isenburg et al. 2001]. They too create geometry from combinatorics by prescribing length 1 to all edges, and solving a nonlinear system for the embedding. Prescribing correct length values to all diagonals in addition is similar to minimizing for ER as in our system. While obtaining similar results (Fig. 23), length optimization fails to achieve the ground-truth (developable) minimum as our optimization does.

Effect of different patterns. In Fig. 24, we test how the geometry changes when the combinatorics of the pattern is altered, to see how predictable the results can be for a designer. We take a quadrilateral mesh, and subdivide it by adding diagonals, while taking care for a more or less even valence of 6 in the result. We choose a blend of $\lambda_{MR} = 1$, $\lambda_{ER} = 0.1$, and increasingly add more positional




Fig. 23. Comparing to connectivity shapes. Left: original mesh with developable-perfect connectivity. Middle: connectivity shapes result. Right: our results reaches the ground-truth minimum.




Fig. 24. Different patterns with similar positional constraints. Left to right: originals, with 4 positional constraints (almost the same spherical results), 8, and 16 constraints. The difference is more apparent along boundaries.

constraints. As expected, the difference grows with the amount of positional constraints, as the mesh is more constrained, and the deviation from regularity of each vertex star manifests somewhat differently in each pattern. Also expectedly, the difference is bigger around boundaries, and specifically where the natural angle sums are different. Nevertheless, both patterns still exhibit “as-spherical-as-possible” behavior, that is comfortable to predict.

Mesh under boundary constraints. The example in Fig. 25 complements the discussion about the behavior of different patterns. The boundary is constrained in both a quad and a diagonal-subdivided triangle mesh. The quad mesh is perfect, and thus its own final result. The natural-boundary perfect solution for the triangle mesh is a $\frac{\pi}{3}$ parallelogram. Constraining the boundary causes an unnatural (though subjectively appealing) result where the inside “bulges” out, in order to create a Möbius-regular mesh. The conclusion of both examples is that the “as-spherical-as-possible” behavior is expected if the mesh is not over-constrained, and otherwise the regularity of the elements takes precedence over cosphericality.

Unconventional patterns. Our algorithm does not decompose patterns, in the sense of extracting directional or repetitive information, as done in [Jiang et al. 2015]. Thus, we could potentially optimize any possible pattern. Three such examples are shown in Fig. 26. We optimize a pentagonal-quad Penrose pattern for Möbius regularity, a pure pentagonal pattern (a subdivided hexagonal mesh), and a rhombic valence 3 mesh for Euclidean regularity. The Penrose pattern, optimized for Möbius regularity, exhibits the same as-spherical-as-possible appearance that we expect, even though the faces are not regular. This is expected since we canonicalize




Fig. 25. Overconstrained optimization. Left to right: quad (original and perfect alike), triangle original, perfect natural-boundary result, and “bulging” result with constrained boundary.




Fig. 26. Unconventional patterns. *Left*: a Penrose pattern designed with Möbius regularity behaves like our other examples. *Center*: a pure pentagonal mesh becomes a mesh of glued Dodecahedrons, that are as Euclidean-regular as possible. *Right*: the rhombic Camel, under Euclidean regularity appears as being made out of small cubes.

the vertex stars. The pentagonal pattern, optimized for Euclidean regularity, exhibits an interesting, perhaps predictable, result: the mesh seems to be made from many glued dodecahedra, reproducing a stock example of Fig. 4. The Camel follows suite: the valence 3 nodes become small cubes, and we get a sort of “Tetrisoid”. It is then inspiring to explore how original patterns would behave under our regularity measures, even without initial expectations. The patterns are made by the subdivision rules in [Akleman et al. 2005].

8 DISCUSSION

Nonlinearity and constraints. The biggest limitation in our work is the nonlinear objective. As such, we can guarantee neither convergence, nor reaching a global optimum, though our results empirically demonstrate robustness. Our algorithm relies on positional constraints to achieve uniqueness. While setting handles is a popular designing endeavor, it requires some experience to set positional constraints that would result in an ideal mesh. Moreover, it is difficult to tell how many constraints are needed for the solution to be unique. Empirically, we found out that setting $g + 3$ positional constraints (where g is the genus) is sufficient.

Limitations. Some patterns admit unsatisfactory solutions, inherent to our energies. We demonstrate two problematic cases in Fig. 27. Möbius transformations contain inversions, potentially producing considerable scale differences in Möbius-regular meshes. This can be alleviated by different positional constraints, and a more dominant component of Euclidean regularity. Conversely, Euclidean-regular




Fig. 27. Fail cases. Left to right: original, optimized for MR with scale differences, and optimized for ER with unordered strips.

triangle meshes may also have an unappealing look in some cases: rather than the intuitive developability that quad meshes exhibit, and while triangles are relatively regular, no developable strips emerge. Identifying such strips, and biasing the optimization accordingly, is likely essential to achieve this effect.

Irregular patterns. Our algorithm currently optimizes meshes under the assumption that the faces are ideally regular. As such, it forces a different symmetry than the one natural to irregular patterns. For instance, for patterns with rhombic faces. This is the reason for the behavior in Fig. 26. Nevertheless, these patterns also have predictable tangent polygons with angles that can be pre-computed. We could thus incorporate such symmetries seamlessly into our method, in a future extension of our work.

Future Work. We would like to apply our work to conformal deformations as well, as discussed in Section 6, preserving surface details relative to a given input mesh, and not necessarily an ideal one. In addition, we would like to explore alternative ways to compute CMC surfaces. Finally, we would like to explore the mathematical properties of regularity with regards to discrete mean curvatures.

ACKNOWLEDGMENTS

We thank Ron VanderFeesten for his help in making the visual aids, and Helmut Pottmann for his major support.

REFERENCES

- Sameer Agarwal, Keir Mierle, and Others. Ceres Solver. <http://ceres-solver.org>.
- Ergun Akleman, Vinod Srinivasan, and Esan Mandal. 2005. Remeshing Schemes for Semi-Regular Tilings. In *Proc. SMI*. IEEE Computer Society, 44–50.
- de Blob. <http://www.designboom.com/architecture/massimiliano-fuksas-de-blob-at-september-18-square/>.
- Alexander I. Bobenko. 2005. A Conformal Energy for Simplicial Surfaces. *Combinatorial and Computational Geom.* (2005), 133–143.
- Alexander I. Bobenko, Tim Hoffmann, and Boris A. Springborn. 2006. Minimal surfaces from circle patterns: geometry from combinatorics. *Ann. of Math.* (2) 164, 1 (2006), 231–264.
- Alexander I. Bobenko, Helmut Pottmann, and Johannes Wallner. 2010. A curvature theory for discrete surfaces based on mesh parallelity. *Math. Ann.* 348, 1 (2010), 1–24.
- Alexander I. Bobenko and Peter Schröder. 2005. Discrete Willmore Flow. In *Proc. SGP*. 101–110.
- Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Levy. 2010. *Polygon Mesh Processing*.
- Sofien Bouaziz, Mario Deuss, Yuliy Schwartzburg, Thibaut Weise, and Mark Pauly. 2012. Shape-Up: Shaping Discrete Geometry with Projections. *Comput. Graph. Forum* 31, 5 (2012), 1657–1667.
- Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Robust Fairing via Conformal Curvature Flow. *ACM Trans. Graph.* 32, 4 (2013).

- Eye. <https://www.eyefilm.nl/>.
- Branko Grünbaum and G C Shephard. 1986. *Tilings and Patterns*. W. H. Freeman & Co., New York, NY, USA.
- Ewain Gwynne and Matvei Libine. 2012. On a Quaternionic Analogue of the Cross-Ratio. *Adv. Appl. Clifford Algebr.* 22, 4 (2012), 1041–1053.
- Udo Hertrich-Jeromin. 2003. *Introduction to Möbius differential geometry*. London Mathematical Society Lecture Note Series, Vol. 300. Cambridge University Press. xii+413 pages.
- Martin Isenburg, Stefan Gumhold, and Craig Gotsman. 2001. Connectivity Shapes. In *Proc. Conf. on Visualization*. IEEE Computer Society, 135–142.
- Caigui Jiang, Chengcheng Tang, Amir Vaxman, Peter Wonka, and Helmut Pottmann. 2015. Polyhedral Patterns. *ACM Trans. Graph.* 34, 6 (2015), 172:1–172:12.
- Hana Kouřimská, Lara Skuppin, and Boris Springborn. 2016. A Variational Principle for Cyclic Polygons with Prescribed Edge Lengths. In *Advances in Discrete Differential Geometry*, Alexander I. Bobenko (Ed.). Springer Berlin Heidelberg, 177–195. KREOD. <http://www.kreod.com/>.
- Yufei Li, Yang Liu, and Wenping Wang. 2015. Planar Hexagonal Meshing for Architecture. *IEEE Trans. on Vis. and Comp. Graph.* 21, 1 (2015), 95–106.
- Christian Müller and Johannes Wallner. 2010. Oriented mixed area and discrete minimal surfaces. *Discrete Comput. Geom.* 43, 2 (2010), 303–320.
- Museo Soumaya. <http://www.soumaya.com.mx/>.
- Jorge Nocedal and Stephen J. Wright. 2006. *Numerical Optimization*. Springer.
- Hao Pan, Yi-King Choi, Yang Liu, Wenchao Hu, Qiang Du, Konrad Polthier, Caiming Zhang, and Wenping Wang. 2012. Robust Modeling of Constant Mean Curvature Surfaces. *ACM Trans. Graph.* 31, 4 (2012), 85:1–85:11.
- Ulrich Pinkall and Konrad Polthier. 1993. Computing discrete minimal surfaces and their conjugates. *Experiment. Math.* 2, 1 (1993), 15–36.
- Helmut Pottmann, Michael Eigensatz, Amir Vaxman, and Johannes Wallner. 2015. Architectural Geometry. *Computers and Graphics* 47 (2015), 145–164.
- Helmut Pottmann, Yang Liu, Johannes Wallner, Alexander Bobenko, and Wenping Wang. 2007. Geometry of Multi-layer Freeform Structures for Architecture. *ACM Trans. Graph.* 26, 3 (2007), 65:1–65:11.
- Rinus Roelofs. <http://www.rinusroelofs.nl/>.
- Boris Springborn, Peter Schröder, and Ulrich Pinkall. 2008. Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* 27, 3 (2008), 1–11.
- Chengcheng Tang, Xiang Sun, Alexandra Gomes, Johannes Wallner, and Helmut Pottmann. 2014. Form-finding with Polyhedral Meshes Made Simple. *ACM Trans. Graph.* 33, 4 (2014), 70:1–70:9.
- Marco Tarini, Kai Hormann, Paolo Cignoni, and Claudio Montani. 2004. PolyCube-Maps. *ACM Trans. Graph.* 23, 3 (2004), 853–860.
- Amir Vaxman. 2016. libhedra: geometric processing and optimization of polygonal meshes. (2016). <https://github.com/avaxman/libhedra>.
- Amir Vaxman, Christian Müller, and Ofir Weber. 2015. Conformal mesh deformations with Möbius transformations. *ACM Trans. Graph.* 34, 4 (2015), 55:1–55:11.

A CORNER TANGENT AND CROSS-RATIO

To prove some technical Lemmas we first consider the following geometric property, which can be easily verified.

LEMMA A.1. *Let $p_i, p_j, p_k \in \mathbb{R}^n$ be three points. Then $(p_i - p_j) \| p_i - p_k \|^2 - (p_i - p_k) \| p_i - p_j \|^2$ is the direction of the tangent of the circumcircle to the triangle $p_i p_j p_k$ at p_i .*

PROOF OF LEMMA 4.1. Note that $q \in \text{Im } \mathbb{H}$ implies $q^{-1} = -q/|q|^2$. Using the definition of the corner tangent yields

$$t[i, j, k] = q_{ij}^{-1} + q_{jk}^{-1} = -q_{ij}/|q_{ij}|^2 - q_{jk}/|q_{jk}|^2.$$

Consequently, Lemma A.1 concludes the proof. \square

LEMMA A.2. *Let $q_i, q_j, q_k, q_l \in \text{Im } \mathbb{H} \cong \mathbb{R}^3$ be four points not lying on a common circle. Then, the imaginary part of the cross-ratio is the normal of the circumsphere (or plane) at q_i . i.e., for a proper circumsphere with center c we have $\text{Im } \text{cr}[i, j, k, l] \parallel (c - q_i)$.*

PROOF. We compute the cross-ratio in terms of corner tangent abbreviated by $t_{f,i} = t[k, i, j]$ and $t_{g,i} = t[k, i, l]$:

$$\begin{aligned} \text{cr}[q_i, q_j, q_k, q_l] &= q_{ji} q_{kj}^{-1} q_{lk} q_{il}^{-1} = q_{ji} q_{kj}^{-1} (q_{ki} q_{ki}^{-1}) q_{lk} q_{il}^{-1} \\ &= (q_{ki}^{-1} q_{kj} q_{ji}^{-1})^{-1} (q_{ki}^{-1} q_{lk} q_{il}^{-1}) \\ &= (q_{ki}^{-1} q_{jk} q_{ij}^{-1})^{-1} (q_{ki}^{-1} q_{lk} q_{il}^{-1}) = t_{f,i}^{-1} t_{g,i}, \end{aligned}$$

where Lemma 4.1 implies that $t_{f,i}^{-1}$ and $t_{g,i}$ are tangent vectors to the circumcircles of triangles $(q_i q_j q_k)$ and $(q_k q_l q_i)$, respectively, both at q_i . Consequently, the cross-ratio is

$$[-\langle t_{f,i}^{-1}, t_{g,i} \rangle, t_{f,i}^{-1} \times t_{g,i}].$$

Its imaginary part is the cross product of tangents to circles on the circumsphere at q_i , hence orthogonal to the circumsphere. \square

The following corollary originally appeared in [Gwynne and Libine 2012, Lemma 15], without the preceding theorem.

COROLLARY A.3. *Let $q_i, q_j, q_k, q_l, q_m \in \text{Im } \mathbb{H} \cong \mathbb{R}^3$ be five points not lying on one circle. Then, these points are cospherical if and only if $\text{Im } \text{cr}[i, j, k, l] \parallel \text{Im } \text{cr}[i, k, l, m]$, i.e., the vector parts of the cross-ratios are parallel (cf. Fig. 10 right).*

B COMPUTATIONS

Corner polygon for mixed meshes. To find the tangent polygon T_v for given lengths (see Section 4.5), we use the following strategy. We assume every face of the 1-ring to be a regular polygon. Assuming unit edge lengths, the length $|t|$ of the corner tangent (or the identical boundary polygon) of a regular d -gon is

$$l_d = |t| = \sqrt{2 - 2 \cos((d-2)\pi/d)}.$$

Obtaining concyclic polygons for prescribed lengths was studied in [Kouřimská et al. 2016]. We take a binary search approach to determine the radius of the circle that we are looking for. For a given radius r , we embed the edges isometrically on the circle of that radius, in cyclic progression, to obtain an *open* concyclic polygon. Consequently, the sum of all sector angles is

$$\sum \arccos(1 - l_n^2/2r^2),$$

and is either larger or smaller than 2π . The polygon is closed (and non-self intersecting) if and only if this angle sum is precisely 2π . This is a monotone function of r , and therefore we can do a binary search until we reach any prescribed precision. Recall that for natural boundary conditions we complete the missing lengths to a closed polygon, and thus this process is not different in this case.