

Hyperbolization of Euclidean Ornaments

Martin von Gagern & Jürgen Richter-Gebert

Lehrstuhl für Geometrie und Visualisierung
Zentrum Mathematik, Technische Universität München, Germany
`gagern@in.tum.de`, `richter@ma.tum.de`

Submitted: Oct 30, 2008; Accepted: May 5, 2009; Published: May 22, 2009

Mathematics Subject Classification: Primary 51M09, 52C26; Secondary 51F15, 53A30, 53A35

Abstract

In this article we outline a method that automatically transforms an Euclidean ornament into a hyperbolic one. The necessary steps are pattern recognition, symmetry detection, extraction of a Euclidean fundamental region, conformal deformation to a hyperbolic fundamental region and tessellation of the hyperbolic plane with this patch. Each of these steps has its own mathematical subtleties that are discussed in this article. In particular, it is discussed which hyperbolic symmetry groups are suitable generalizations of Euclidean wallpaper groups. Furthermore it is shown how one can take advantage of methods from discrete differential geometry in order to perform the conformal deformation of the fundamental region. Finally it is demonstrated how a *reverse pixel lookup* strategy can be used to obtain hyperbolic images with optimal resolution.

1 Introduction

Ornaments and regular tiling patterns have a long tradition in human culture. Their origins reach back to ancient cultures like China, Egypt, Greece and the Islam. Each of these cultures has its own specific way to create symmetric decorative patterns. Mathematically the underlying structures are well understood. Taking only the pattern into account (and neglecting possible color symmetries) there are just 24 sporadic groups and 2 infinite classes of groups governing the symmetry structure of plane images. The two infinite classes belong to rosette patterns with an n -fold rotational symmetry (having reflections or not). Furthermore there are 7 frieze groups describing patterns that have translational symmetry in a single direction. And last but not least there are the most interesting ones, the 17 wallpaper groups for patterns that admit translational symmetries in at least two independent directions. The frieze and wallpaper groups are infinite symmetry groups.

Structurally we have to tell apart at least three different layers when speaking about symmetric patterns – two of them belong to mathematics and one belongs to art. Starting




Figure 1: A planar Euclidean ornament with its underlying symmetry structure

with a concrete symmetric pattern $\mathcal{P}$ we are interested in its symmetry structure. For this we embed the pattern in $\mathbb{R}^2$ and study all Euclidean symmetries of this pattern. One formal way to do this is to define a mapping $\mathcal{P} : \mathbb{R}^2 \rightarrow C$ where C is a suitably rich set that resembles a color space. For instance we could take $C = [0, 1]^3$ to be the red/green/blue color values of each point in the plane. We are looking for all Euclidean transformations that leave the pattern invariant. Let $\mathbf{euc}$ be the set of all possible Euclidean isometries. The symmetry group of the pattern $\mathcal{P}$ is $\text{sym}(\mathcal{P}) := \{g \in \mathbf{euc} \mid \mathcal{P} \circ g = \mathcal{P}\}$, where $\mathcal{P} \circ g$ is a shorthand for first applying the Euclidean symmetry and then looking for the color at the transformed position. The symmetry group is the part of the automorphism group of $\mathcal{P}$ that belongs to $\mathbf{euc}$. In other words we arrive at the same color for every location in the orbit $\text{orb}(x) := \{g(x) \mid g \in \text{sym}(\mathcal{P})\}$. This agrees with Hermann Weyl's famous definition of symmetry: "an object is symmetric if it remains the same under some transformations". Under the assumption that $\mathcal{P}$ does not admit continuous symmetries, each planar pattern $\mathcal{P}$ has a symmetry group falling into one of the (conjugacy) classes mentioned above. (The extreme case of a pattern without any repetitions is covered by the rosette group that has only a 1-fold rotation.) The conjugacy class of the symmetry group is in a sense the highest level of abstraction.

On a less abstract level each pattern is associated to a *concrete* symmetry group. Within the same conjugacy class these concrete groups may still differ by a variety of parameters, like scaling, rotations, etc. Factoring out Euclidean transformations, some of these classes turn out to have just one unique representation. This happens for all rosette and frieze groups. For wallpaper groups this effect typically arises when a rotational symmetry of order at least three is present. We call such groups *highly symmetric*. The remaining wallpaper groups still have one or two degrees of freedom in their concrete geometric representation. They correspond to an anisotropic stretching in a certain direction or an angle between the generating directions. The specific geometric representation of a group is the second mathematical level we have to consider.

Finally, there is the artistic level, which is responsible for the concrete motive that is repeated in the ornament (see Figure 1). Here one has all artistic freedom from complete arbitrariness and noise to sophisticated ornamental designs as they occur for instance in




Figure 2: A tessellation of the hyperbolic plane and M.C. Escher's Circle Limit III, woodcut 1959.

the Islamic ornamental art. The visual appearance of an ornament may vary drastically with its artistic input. There is a clear concept that helps to draw a border between the artistic and the mathematical level of an ornamental pattern: the *fundamental domain*. This is a (connected) region in the original pattern that contains exactly one point from each orbit of the pattern. Thus within (the interior of) a fundamental domain there is no repetition at all and one has full artistic freedom. If each point of one fundamental domain is colored then all the rest of the ornament is determined by the underlying symmetry group. Multiple copies of the fundamental domain generate the pattern. One of the aims of this paper is to provide a method that allows one to take the high-quality artistic input of classical ornaments and turn them into patterns with new and amazing kinds of symmetry. In other words we are interested in the generation of artistically valuable fundamental domains for interesting symmetry groups.

As mentioned above the structure of Euclidean ornaments is well understood and in a sense a classical topic of geometry (see [9, 10, 18, 19]). It very seldom happens that there are specific moments that represent a pivot in the interaction of mathematics and art. But in the case of ornaments there seems to have been at least one such. At a conference in 1954 the famous Dutch artist and lithograph M. C. Escher and H.S.M. Coxeter first met. They exchanged many ideas on the interrelations of mathematics, geometry and art. The conference resulted in a friendship and an exchange of letters. One of these letters of Coxeter to Escher contained a rendering of a hyperbolic tessellation in the Poincaré disk as it could also be found in Felix Klein's and Poincaré's work (see Figure 2 left). Escher claimed that he was *shocked* by this picture since it resolved a problem he had been struggling with for several years: *How to fit infinitely many similar objects into the finite limits of a circle*. Inspired by this drawing Escher produced a sequence of pictures – his famous *Circle Limit I – IV* (see Figure 2 right). And in reaction to this Coxeter wrote a beautiful article describing the subtle mathematics of these pictures [5, 6].

The hyperbolic plane allows several (hyperbolic) rotation centers of arbitrary degree.




Figure 3: The ultimate aim of this article: Transforming a scanned Euclidean input image to a perfect hyperbolic ornament

Hence in contrast to the Euclidean plane in the hyperbolic plane there are infinitely many structurally different symmetry groups. Nevertheless the amount of available artistically interesting images of *hyperbolic ornaments* is very limited, and still the few examples created by Escher are the outstanding period of this art. The reasons for this are obvious. A person creating hyperbolic ornaments has to cope with several problems and to be equipped with several skills. He/she has to be artistically gifted and has to know the mathematical backgrounds of hyperbolic tessellations including the highly non-trivial geometric construction principles. Furthermore he/she has to find a way to produce infinitely many ever so tiny copies of the object that has to be repeated (the fundamental region). There are several attempts to create such images by computer. Some of them [3] try to automatically create fundamental regions that match a given concrete shape. Others try to manually transform Euclidean patterns to create similar hyperbolic patterns [8]. There are also attempts [16] to take original picture parts of Escher images, deform them manually to fit approximately to a hyperbolic fundamental region and use this as a basis for a tessellation. These latter come in a sense closest to the ones generated by the methods described in this article. However, they suffer from visual seams which appear at the boundaries of the patches.

Also there are several programs that can be used to create individual designs for a specific symmetry group by mouse interactions (drawing lines, placing objects). Most of these programs also suffer from the problem of creating infinitely many objects in finite time and are of doubtful artistic quality.

In this article we want to take a different approach. We want to outline a method that allows one to take a classical ornamental pattern, feed it into a computer and create corresponding seamless hyperbolic patterns. Thus we want to keep the artistic content (of the old masters), while changing the symmetry structure of the overall picture. Figure 3 exemplifies input and a possible output of the program. The present article outlines the

different steps of this production process.

- **Pattern recognition:** We use autocorrelation methods based on discrete Fourier transforms to detect symmetries of the original picture and to identify the structure of its symmetry group.
- **Finding corresponding hyperbolic groups:** An analysis of the symmetry group is used to determine which hyperbolic groups match the original symmetry group.
- **Conformal deformation:** We create a conformal deformation of the original fundamental region to the corresponding hyperbolic region. The theory of this step is governed by Riemann's famous mapping theorem. However, practically feasible implementations of this step need sophisticated methods from the relatively new field of discrete differential geometry [2]. In particular, we use discrete conformal mappings as introduced in [20].
- **Creating the image:** Finally, the unit Poincaré disk has to be covered with infinitely many copies of the fundamental region according to the chosen symmetry group. In order to obtain a perfect raster image we use a pixel oriented algorithm that calculates for each pixel its corresponding color.

The rest of the article is organized as follows. After outlining some elementary concepts and definitions we focus in Section 3 on the task of transforming a Euclidean ornament into a hyperbolic one. Section 4 is dedicated to the pre- and post-processing – this means the pattern recognition and the task of filling the entire Poincaré disc with the pattern. In Section 5 we present a collection of interesting examples. Finally, Section 6 points to further projects and problems in this context that are not covered by this article. The article is meant as an overview of the subjects and our methods and omits several technical details in order to emphasize the overall picture.

2 Concepts

This section covers several fundamental concepts needed for the rest of the article. We assume that the reader is familiar with elementary concepts of wallpaper groups and with hyperbolic geometry.

2.1 The Crystallographic Groups

The input pictures that we will use will all be Euclidean wallpaper ornaments. These pictures admit translational symmetries in at least two independent directions. Besides translations also rotations, reflections and glide reflections are allowed. The classification of these groups dates back to the late 19th century and is due to Fedorov, Schönflies and Barlow [1, 9, 10, 18, 19]. There are exactly 17 conjugacy types of such groups. Different nomenclatures have been proposed during the centuries. We here will stick to the IUC




Table 1: Ornaments and their wallpaper groups

notation, which is the notation for the symmetry groups adopted by the International Union of Crystallography in 1952 [14]. A list of all 17 groups is given in Table 1 (the table first appeared in [11] and is based on the beautiful collection of classical ornaments published by Owen Jones in 1910 [15]). The table gives one example for each of the groups and overlays a diagram indicating the symmetry structure of the group. Double lines represent reflections. Polygons represent rotation centers and dashed lines represent glide reflections. In the nomenclature a letter m indicates the presence of a reflection, a g indicates a glide reflection and a number $n > 1$ indicates the presence of an n -fold rotation.

A concept that will turn out important for our purposes will be the *fundamental region* (or *fundamental domain*) F of a wallpaper group $G \subset \mathbf{euc}$. The orbit under G of a point p is the equivalence class $\{g \circ p \mid g \in G\}$.

Definition 1 A fundamental region F of a wallpaper group $G \subset \mathbf{euc}$ is a connected and closed region in $\mathbb{R}^2$ that contains at least one point from every orbit of the group, such that the images of F under G cover $\mathbb{R}^2$ and such that two such images have no interior points in common.

Fundamental regions with fractal boundary are objects of current research. However, we will neglect this case here. For Euclidean wallpaper groups we will only consider polygonal fundamental regions. Later on in the hyperbolic case we will also admit hyperbolic polygons (i.e. bounded by Euclidean circular arcs in the Poincaré disk). While all interior points of F belong to different orbits, it may happen that several boundary points of F belong to the same orbit. For a particular colored symmetric pattern $\mathcal{P} : \mathbb{R}^2 \rightarrow C$ the artistic content is given by the behavior within a fundamental region. The whole pattern is generated by gluing infinitely many copies (some of them perhaps reflected) of the fundamental region along their boundaries.

Some of the wallpaper groups turn out to be two-dimensional reflection groups or are closely related to them. The groups p4m, p3m1 and p6m correspond to the three Euclidean triangular kaleidoscopic fundamental regions with corner angles $(90^\circ, 45^\circ, 45^\circ)$, $(60^\circ, 60^\circ, 60^\circ)$ and $(90^\circ, 60^\circ, 30^\circ)$, respectively. The groups p4, p3 and p6 are the corresponding reflection free subgroups. These groups are index two subgroups of the corresponding reflection groups. The groups p31m and p4g can also be considered as index two subgroups of p6m and p4m, respectively. All these groups have only one possible geometric conjugacy class, since they are (up to a global Euclidean transformation) uniquely determined by the shape of the underlying triangles. There is one more reflection group, namely pmm. It corresponds to a rectangular kaleidoscope with four 90° corners. It has a one parameter family of conjugacy classes parametrized by the ratio of two adjacent edges.

Here we will discuss “hyperbolizations” of wallpaper groups that contain at least one proper center of rotation. In addition to the above mentioned groups these are p2, pmg, cmm, pgg. The groups pmg, cmm and pgg have a one parameter family of geometric conjugacy classes. The group p2 has two geometric parameters (its fundamental region can be an arbitrary triangle).

2.2 Hyperbolic Ornaments and Hyperbolization

Analogously to Euclidean wallpaper groups we will also consider hyperbolic ones. By default we will use the Poincaré model of the hyperbolic plane. The hyperbolic plane is the interior of the unit disk $\mathbb{H} := \{z \in \mathbb{C} \mid |z| < 1\}$ in $\mathbb{C}$ equipped with a certain metric (see for instance [12] for further details). In this model hyperbolic lines correspond to Euclidean circular arcs that intersect the boundary of the Poincaré disc orthogonally. The model is conformal in the sense that hyperbolic angles between hyperbolic lines correspond to the intersection angle of the corresponding circles. Orientation preserving hyperbolic isometries are given by Möbius transformations $\mu : \mathbb{H} \rightarrow \mathbb{H}$ that leave the unit circle as a whole invariant. Orientation reversing hyperbolic isometries are obtained by combining a Möbius transformation with complex conjugation. By **hyp** we denote the set of all these isometries. Now a hyperbolic ornament is a color pattern $\mathcal{P} : \mathbb{H} \rightarrow C$ that admits a discrete symmetry group $\text{sym}(\mathcal{P}) := \{g \in \mathbf{hyp} \mid \mathcal{P} \circ g = \mathcal{P}\}$. In contrast to the Euclidean case in the hyperbolic plane there are infinitely many different discrete symmetry groups having more than one center of rotation.

The simplest of these groups are the triangular reflection groups generated by triangular kaleidoscopes with angles $\frac{\pi}{k}, \frac{\pi}{m}, \frac{\pi}{n}$, where $k, m, n \in \mathbb{N}$ and $\frac{1}{k} + \frac{1}{m} + \frac{1}{n} < 1$. Already in this case there are infinitely many different such groups governed by the different choices of k, m, n . In our approach these groups will play a special role, since they are especially easy to deal with.

Before coming back to these groups we will define what we mean by a *hyperbolization* of a Euclidean ornament. As mentioned in the introduction we want to maintain as much as possible of the artistic input. In particular shapes and angles should only be minimally distorted by the deformation of the fundamental cell. For this reason we require the deformation to be a conformal mapping except for the rotation centers. We identify the Euclidean plane with the complex numbers $\mathbb{C}$. We will define hyperbolization by an analytic continuation process of a function f that is consistent with the symmetry of the pattern $\mathcal{P}$.

Definition 2 *Given a Euclidean pattern $\mathcal{P}$ that belongs to a specific wallpaper group $G \subset \mathbf{euc}$. Furthermore, let $\mathcal{P}' : \mathbb{H} \rightarrow C$ be a hyperbolic ornament with some underlying discrete hyperbolic group $G' \subset \mathbf{hyp}$. Let $U \subset \mathbb{C}$ be a small region that does not contain a rotation center of G . We say that $\mathcal{P}'$ is a hyperbolization of $\mathcal{P}$ if there is a conformal map $f : U \rightarrow \mathbb{H}$ such that*

- $\mathcal{P}(z) = \mathcal{P}'(f(z))$ for all $z \in U$ and
- For any analytic path $\psi : [0, 1] \rightarrow \mathbb{C}$ with $\psi(0) \in U$ that avoids all rotation centers of G the analytic continuation $\tilde{f}(z)$ of f satisfies $\mathcal{P}(\psi(1)) = \mathcal{P}'(\tilde{f}(\psi(1)))$.

The analytic continuation of f connects (conformally) the color patterns $\mathcal{P}$ and $\mathcal{P}'$. If we trace any path in $\mathbb{C}$ that avoids rotation centers the colors of the two patterns will coincide for z and $\tilde{f}(z)$, respectively. It should be mentioned that at the rotation centers considerable monodromy could occur. Imagine that we have a fourfold rotation center r

in $\mathcal{P}$ that gets mapped to a fivefold rotation center r' in $\mathcal{P}'$. If we start with some point z sufficiently close to r and with a corresponding image point $\tilde{f}(z)$ then each full cycle around r causes the image point to perform a $\frac{4}{5}$ turn around r' . Only after five full cycles the image z and $\tilde{f}(z)$ are both back to their original position.

2.3 Orbifolds

One of the most successful concepts in the study of symmetric patterns is the underlying *orbifold*. In our context it will sometimes be useful to perform considerations on the level of orbifolds rather than on the level of fundamental regions. Similar to wallpaper groups the orbifold can also be considered on a geometric and on a combinatorial (topological) level. The orbifold is an abstract two-dimensional manifold, that contains exactly one point from each orbit of a wallpaper group. For a detailed introduction to the theory of orbifolds we refer to the beautiful book [4]. Here we only introduce the concepts necessary in our context. In the Euclidean plane we consider a polygonal fundamental region F of a wallpaper group G . The orbifold now can be considered as a copy of a single fundamental cell where boundary points that are identified via G are glued together. The edges of F that belong to reflections become the boundary of the orbifold, while centers of rotation become corner points on the boundary or cone points in the interior of the orbifold, depending on whether they lie on an axis of reflection or not. Figure 4 illustrates the concept. It shows the famous Angel and Devil ornament of M. C. Escher. The symmetry group of this picture is $p4g$. The drawing in the middle highlights a fundamental region. Cutting out this right-angled triangle and identifying the two edges labeled a we obtain the orbifold. In this case it has one cone point of order 4. The hypotenuse of the triangle corresponds to a mirror line and becomes the boundary of the orbifold.

Our approach to hyperbolization will take the orbifold of a wallpaper group and attempt to change the order of the corners and cone points. Thus in the example of the Angel and Devil picture above we might be interested to produce a version of the picture in which for instance five instead of four tips of the wings fit around the cone point. This is no longer possible in the Euclidean plane. However one can find a suitable embedding in the hyperbolic plane. A corresponding picture is shown on the right of Figure 4 together with the deformed fundamental region. Notice that this region is no longer bounded by straight lines but by circular arcs (which are hyperbolic straight lines). In the example above the conformal relation between the Euclidean and the hyperbolic picture is already induced by a suitable conformal mapping of the fundamental region.

3 Computing Hyperbolizations

The following sections deal with various aspects of hyperbolization. We will outline theoretical as well as computational approaches. In particular, we will describe cases that are particularly easy (triangle groups and their relatives) and exemplify difficulties for the other groups.




Figure 4: The fundamental region of an ornament can be glued along parts of the boundary to become an orbifold. Deformed orbifolds can again be used as fundamental regions for hyperbolization. (Pictures based on M.C. Escher, symmetry drawing number 45, 1941.)

3.1 Triangle Groups and their Relatives

The restriction of conformality is a relatively strong requirement, but in the case of triangle groups it is reasonably simple to fulfill. We will first outline the case of hyperbolization of reflection groups generated by a triangular kaleidoscope.

Assume that the pattern to be transformed belongs to one of the triangular reflection groups $p4m$, $p3m1$ or $p6m$. They have a triangular fundamental cell Δ_{euc} . The only reasonable hyperbolizations will have a fundamental region of comparable shape and reflection behavior. Thus a hyperbolization belongs to a hyperbolic triangular group with corner angles $\frac{\pi}{k}$, $\frac{\pi}{m}$ and $\frac{\pi}{n}$ where $\frac{1}{k} + \frac{1}{m} + \frac{1}{n} < 1$. The corresponding fundamental region is a hyperbolic triangle Δ_{hyp} . In order to be conformal we need a mapping that maps Δ_{euc} to Δ_{hyp} and is conformal except for the corner points. At least theoretically such a mapping is simple to construct via the Riemann mapping theorem. By this there exists a conformal map f_1 from the interior of Δ_{euc} to the upper half-plane $\mathbb{C}^+$. This map is unique up to a Möbius transformation. The only regions where this map is not conformal is at the corners of the triangle. There it behaves like a function z^α . If we fix the positions of the image's corners on the real axis the mapping is uniquely determined. In a similar way we can define a mapping f_2 from Δ_{hyp} to $\mathbb{C}^+$ (for this we chose w.l.o.g. the same images of corner points on the real axis). Both maps and their inverses are conformal. The desired map from Δ_{euc} to Δ_{hyp} now is given by $(f_2^{-1}) \circ f_1$. The map f_1 is a special case of a Schwarz-Christoffel mapping and f_2 is a variant of Schwarz-Christoffel mappings for circular arcs (see [7]).

At least theoretically there is an explicit way to express f_1 and f_2 . This can be done by the use of the hypergeometric function ${}_2F_1(a, b, c, z)$. It goes back to Schwarz's original work in which these maps were first introduced particularly in the context of

covering spaces of hyperbolic reflection groups. If we want to map the upper half-plane to a triangle with corner angles α, β, γ this can be done by the formula

$$\phi(z) = \frac{z^{1-c} {}_2F_1(a', b', c', z)}{{}_2F_1(a, b, c, z)}$$

with $a = (1 - \alpha + \beta - \gamma)/2$, $b = (1 - \alpha - \beta - \gamma)/2$, $c = 1 - \alpha$, $a' = a - c + 1$, $b' = b - c + 1$ and $c' = 2 - c$ (see [13]). This formula maps the real axis to a triangular region with the required corner angles. The preimages of the corner points are 0, 1 and ∞ . Calculating the inverse of this map requires sophisticated numerical methods since the map $\phi(z)$ may become highly degenerate.

Our first attempts to hyperbolization were based on this approach. Unfortunately, these maps are difficult to calculate and the known numerical methods (see for instance [7]) become unstable for circular arc polygons, due to so-called “crowding effects”. For these reasons our first attempts at a concrete implementation took about one hour CPU time for the deformation process, even if only tiny angle deformations were required. In the meantime we switched to more sophisticated methods from discrete differential geometry. We will come back to this issue in Section 3.3.

Having transformed a single triangle the rest of the hyperbolic unit disk $\mathbb{H}$ can be filled consistently by the use of the *Schwarz Reflection Principle* (SRP). In its basic form this principle can be stated most easily within the framework of the complex plane. If $G \subset \mathbb{C}$ is some region in the upper half-plane partially bounded by a real line segment S and $f : G \rightarrow \mathbb{C}$ is a conformal map that maps S again onto the reals axis then f can be analytically and conformally extended to a region $\overline{G} := \{\overline{g} \mid g \in G\}$ by the function $f'(z) := \overline{f(\overline{z})}$. To be more specific: If we have a triangle Δ in the upper half-plane with one side S on the real axis and this triangle is mapped to some region $f(\Delta)$ in a way that $f(S)$ is again a real segment then this map can be extended as follows: The mirror image of Δ along the real axes is mapped to the mirror image of $f(\Delta)$ according to $\overline{f(\overline{z})}$. The map will be conformal also along the real axis. By suitable conjugation with Möbius transformations the SRP can also be applied to regions bounded by arbitrary circular arcs (of possibly infinite radius). By this in a hyperbolization of a reflection group the triangular image of one triangle determines the images of all three neighboring triangles. If the corner angles of the image triangle are divisors of π then the iterated continuation of the map closes up nicely around the new rotation centers (in fact we get a very interesting monodromy behavior of the map around the rotation centers, however we will neglect this issue here).

One should notice that the triangular reflection groups are easy to handle for the following reason: The Schwarz Christoffel map that relates the upper half-plane to the interior of a (circular arc) polygon is unique up to Möbius transformation. Thus in the case of triangles this imposes no serious restriction on the position of preimages of the triangle corners. They can be entirely controlled by the Möbius transformation. The situation changes if the fundamental region has more than three corners. However, the fact that p3, p4, p6, p31m and p4g are closely related to triangle groups helps us, together




Figure 5: Families of triangle-based orbifolds

with the SRP, to deal with these groups easily as well. Here is a census of the fundamental regions of these groups, also illustrated in Figure 5.

Reflected Triangle (p6m, p3m1, p4m): The fundamental cell of the ornament is a triangle which is bounded by reflections on all three sides. The interior angles of $\frac{\pi}{k}$, $\frac{\pi}{m}$ and $\frac{\pi}{n}$ correspond directly to centers of rotation of order resp. k , m and n , which are implied by the reflections. Any integral values that fulfill the angle sum inequality for a hyperbolic triangle are valid.

Rotated Kite (p6, p3, p4): The fundamental cell of the ornament can be chosen as a deltoid with equal angles $\frac{\pi}{k}$ symmetric to one another and angles $\frac{2\pi}{m}$ and $\frac{2\pi}{n}$ on the axis of rotation. Again they correspond to centers of rotation of order resp. k , m and n . The two symmetric rotations of order k belong to a single transitivity class. The angle sum inequality has to be fulfilled for one half of the kite.

Reflected Kite (p31m, p4g): The fundamental cell of the ornament can be chosen as a deltoid with reflections along two symmetric sides. The symmetric angle has to be a right angle, while the angles on the axis of symmetry are $\frac{\pi}{m}$ at the corner with the reflections and $\frac{2\pi}{n}$ at the corner without reflections, corresponding to centers of rotation of order resp. m and n . The angle sum inequality has to be fulfilled for one half of the kite.

Thus the fundamental cells of the groups p3, p4, p6, p31m, p4g consist of two copies of the triangles for p3m1, p4m or p6m. Hence, if we already have a conformal map on this triangle it can be easily extended to the whole kite by the SRP. By this we also get a mapping of the quadrangular fundamental regions in these cases. Figure 6 shows the deformation of a fundamental region of an ornamental pattern with p4 as underlying wallpaper group.




Figure 6: Conformal deformation of a Euclidean fundamental region to a hyperbolic one.

Multiple hyperbolic copies of this patch will perfectly cover the hyperbolic plane, as Figure 7 shows. Here our standard example which belongs to a $p4$ has been transformed to a suitable hyperbolic group. The original fundamental region is a square. One of the 90° angles (the one in the center of the flower) has been altered to an angle of 60° . The one centered in the little square as been altered to 72° . It is instructive to observe how different parts of the picture change their appearance. The yellow wheel around the flower is now connected to six other wheels, while the little square became a pentagon.




Figure 7: Filling the unit disc with hyperbolic copies of a patch.

3.2 Parametric Symmetry Groups

The situation becomes by far more subtle for the groups with lower symmetry. This section deals with the specific problems that arise if the shape of the fundamental region is not already determined by the combinatorics of the symmetry group.

3.2.1 Pmm

We first consider the case of the group pmm . The group pmm has a rectangular fundamental region, with reflections along all boundary edges. We now want to conformally deform this rectangle to obtain a suitable hyperbolic quadrangle with altered angles at the corners. In order to cover the hyperbolic plane properly, every corner angle must be a divisor of π . From a combinatorial point of view, this is the only restriction. Thus we can set the angles to, say, 60° at each corner and obtain infinitely many possible tessellations of $\mathbb{H}$. In fact, the hyperbolic length of one of the sides together with the corner angles uniquely determines the shape of the tile. Figure 8 illustrates a sequence of such tessellations coming from a continuous change of the length parameter.

Our definition of a proper hyperbolization however requires that we can *conformally* map the original rectangle in the Euclidean plane to its hyperbolic counterpart. This turns out to be possible for just one specific shape of the hyperbolic tile in our infinite class. The reason for this is as follows: The Riemann mapping theorem tells us that there is a conformal map from the interior of the Euclidean rectangle to the upper half-plane. This map is unique up to Möbius transformations. Hence the position of the images of three of the corner points on the real axis uniquely determines the position of the fourth corner. Similarly, there is such a map for the hyperbolic fundamental region. In order to map corners of the Euclidean tile to corners of the hyperbolic tile the images of the corresponding corner points on the real axis must agree. Three points can be identified via the Möbius transformation. Fitting the fourth point is in fact a condition. There is a notion in conformal mapping theory that singles out this condition – the *conformal modulus* (see [7]). Given a quadrangle (perhaps bounded by circular arcs) we consider




Figure 8: Instances of a parametric hyperbolic symmetry group.

a conformal map to the upper half-plane. The boundary of the quadrangle is mapped to the real axis. The map is unique up to Möbius transformations and the cross ratio of the images of the four corner points is an invariant of this mapping process. Thus two quadrangles can only be conformally mapped to each other if their conformal moduli agree.

It is instructive to study why there was no problem with quadrangular fundamental regions in the case of the symmetry groups $p4$, $p3$, $p6$, $p4g$ and $p31m$. In all these cases the fundamental regions are kites. The fundamental regions of the hyperbolizations are hyperbolic kites as well. All kites have the same conformal modulus -1 since they are conformally equivalent to a square. Hence the shape of the fundamental regions is no restriction to conformality.

In the standard setup of numerical Schwarz-Christoffel mappings (these are conformal mappings of polygonal regions to the upper half-plane) the computation is usually divided into two distinct numerical steps (see [7]). When a polygon is mapped then in a first step the position of the images of the corner points on the real axis is numerically calculated. A second numerical step computes the actual mapping. We will later on see how we circumvent these issues by directly generating a picture of the conformally mapped fundamental region.

3.2.2 $P2$, Pmg , Pgg and Cmm

The situation becomes even worse for the groups $p2$, cmm , pgg and pmg . In these cases we have parametric symmetry groups for which parts of the boundary of the fundamental region have to be identified with other parts of the boundary. There the situation is as follows: The fundamental region of the Euclidean group may be chosen to be a polygon. However, the specific position of these boundary edges of the fundamental region has no intrinsic geometric meaning. A priori there is no reason why such a polygonal region should in a suitable hyperbolization be mapped to a hyperbolic polygonal region. (The situation is different for triangle groups and their relatives. There the kite shape together with the SRP ensures that a polygonal Euclidean fundamental region is mapped to a hyperbolic polygon.)

The above considerations indicate that it is desirable to have a concept of conformal deformation that is not entirely based on the concept of deforming fundamental regions, since their edges are sort of artificial. In Section 3.3.3 we will come back to this issue and describe how approaches from discrete differential geometry can provide methods that work directly on the more appropriate structure of the underlying orbifold.

3.3 Discrete Conformal Deformation

3.3.1 Discrete Conformality

Discrete differential geometry allows for a more feasible alternative to the time-consuming continuous calculations using hypergeometric functions outlined in Section 3.1. Spring-

born et al. recently introduced the concept of discrete conformal equivalence of triangle meshes [20]. They also describe an algorithm to calculate such meshes.

Taken as a black box, the algorithm has the following interface.

Input: a combinatoric description of a triangulation, including

- for each edge its length in the original mesh and
- for each vertex its target angle sum in the result mesh.

Output: a factor f_v for every vertex v such that when all edges are multiplied with the factors of both their endpoints, the resulting mesh fulfills the desired angle sums and is called conformally equivalent.




Figure 9: Example of conformal equivalence

An example of this process is given in Figure 9. To get a flat polygon as the result of such a transformation, all interior vertices should be assigned an angle sum of 2π , and vertices on the boundary π , with the exception of the corners of the polygon.

To transform the interior of each triangle, some kind of interpolation is required. It is possible to use the factors f_v to calculate a projective transformation for each triangle which will not only map corners to corners but also preserve the circumcircle of these corners and guarantee continuity along the triangle edges. This is a nice property of discrete conformality defined in this way. A mapping is discretely conformal if and only if the circumcircle preserving projective maps of adjacent triangles agree on joint edges [20].

3.3.2 High Symmetry Case

With the algorithm for discrete conformal equivalence in place, one can use it to transform a single fundamental domain. To do so, we start with a hyperbolic fundamental domain and transform it to a straight-edged Euclidean fundamental domain with given corner angles. In rough steps the procedure goes as follows:




Figure 10: Discrete conformal transformation of a triangulated fundamental domain

1. Embed the hyperbolic fundamental domain into the Euclidean plane using Poincaré's disk model.
2. Approximate this arc-bounded shape using a sufficiently fine mesh of Euclidean triangles.
3. Designate target angles for the centers of rotation according to the order of these rotations in the original Euclidean symmetry group.
4. Default all other vertices to π along the boundary resp. 2π in the interior of the mesh.
5. Calculate a conformally equivalent mesh for these angles.
6. Adjust the size and position of the transformed mesh to match the original Euclidean fundamental domain.
7. Map interiors of all triangles using circumcircle-preserving projective transformations.

The above algorithm assumes that you already know the shape of the appropriate hyperbolic fundamental domain. This is the case for highly symmetric Euclidean groups, i.e. groups containing a rotation of order at least three. In these cases, the shape of the fundamental domain is already fully determined by the combinatorics of the hyperbolic symmetry group, as was described in Section 3.1.

Using a triangle reflection group as the basis for a hyperbolic symmetry group is a powerful and intuitive concept. In another application written by the first author, a user can define hyperbolic symmetry groups by identifying triangles of a triangular tiling. The transformations defined in this way generate a subgroup of the reflection group. This symmetry group can then be used to draw new hyperbolic ornaments.

3.3.3 Low Symmetry Case

In the less symmetric groups, more complicated issues occur. As explained in Section 3.2, in these cases the corners of the fundamental domain can't be deduced from the combinatorics of the hyperbolic symmetry group, and in general you can choose either the Euclidean or the hyperbolic fundamental domain to be bounded by straight lines, but not both at the same time, as illustrated in Figure 11. To deal with these issues, one can focus on the orbifold instead of a flat fundamental domain. In contrast to the latter, the orbifold contains no arbitrarily chosen boundary edges. The only possible boundary is introduced by reflections, which have to map to reflections and thus remain straight-edged. All other boundaries of the fundamental domain are glued together and thus lose their distinguished role.

One can easily build a mesh representation of the orbifold of the original Euclidean ornament. This is done by triangulating the Euclidean fundamental domain in any suitable manner and then identifying boundary objects according to the topology of the Euclidean group. As a result, each triangle has a fixed position in the fundamental domain, but edges and vertices don't, and triangles not adjacent in the fundamental domain may be adjacent in the orbifold mesh. The resulting mesh in general can't be layed out in the plane. Luckily the input for the calculation of conformally equivalent triangle meshes requires only edge lengths, not vertex positions. The lengths of these edges are well defined and easy to calculate.

Now that we have a mesh of the Euclidean orbifold, we can hyperbolize it. To do so we simply assign different target angles to the corners and cone points of the orbifold, fulfilling the hyperbolic angle sum constraints. There is a hyperbolic version of the conformality algorithm, which uses hyperbolic edge length and angle calculations. Using this, we end up with a discretely conformal hyperbolic orbifold.

To draw the hyperbolic ornament, one has to unroll this orbifold. Starting with an arbitrary triangle and placing it at an arbitrary position, one can repeatedly place adjacent triangles next to those already layed out. For every new triangle to be placed, the position of an edge already placed together with the lengths of all three edges defines the position of the third vertex. In theory, this process could be repeated infinitely in order to cover the whole hyperbolic plane. In practice, this process by itself would be infeasible to draw the whole image. Instead, it suffices to draw a single hyperbolic fundamental domain,




Figure 11: Fundamental domains connecting centers of rotation using hyperbolic straight edges (red) or the images of Euclidean straight edges (blue).

and continue from there using the process described in Section 4.2. Once the layout process has laid out enough centers of rotation, they can be connected using straight hyperbolic lines to form a hyperbolic polygonal fundamental domain. Triangles outside this fundamental domain needn't be considered. With this restriction in place, most triangles of the orbifold result in only one triangle in the interior of the fundamental domain. Only a few triangles of the orbifold result in multiple copies being positioned near different edges of the fundamental domain belonging to the same orbit.

4 Pre- and Post-processing

While the conformal deformation of the fundamental cell is the core of the hyperbolization process, this by itself isn't enough to turn an image of an Euclidean ornament into a complete hyperbolic ornament. There are two more necessary steps. Both of them have their own mathematical subtleties. First of all the deformation algorithm needs proper input data. Hence, before the deformation can even be started, the input image usually has to be made perfectly symmetric, and its symmetry structure has to be analyzed. This data is fed to the deformation algorithm. As described in Section 3 this algorithm produces a suitable hyperbolic fundamental region. In a kind of post-processing this data is used to fill the entire Poincaré disk according to the symmetry structure of the underlying hyperbolic group. The following two sections deal with these two steps.

4.1 Recognition of Ornaments

The input to our process should be an image of any symmetric ornament. In general, such an input will be imperfect: although a human will readily perceive it as being symmetric, it actually isn't strictly so.

Take Figure 12(a) as an example. Although you can tell that the artist intended there to be centers of four-fold rotation in the middle of each green square, you can also have a closer look and see that in fact the color plates are misaligned, and also that the squares do not all look exactly the same. Therefore, strictly speaking the ornament is not symmetric. If one were to use this imperfect image directly as the source of the hyperbolic ornament, seams would be visible in the result. Consider for example the misaligned colors in the center of the square. As the hyperbolization process will turn this four-fold rotation into a five-fold one, it has to be split at some line, where at least one more segment can be inserted. Practically, one would use only a single fundamental domain, which has to line up with itself.

Therefore, as the initial step towards a hyperbolic ornament, we first have to get a perfectly symmetric Euclidean ornament. To help with this, the computer has to analyze the structure of the ornament, i.e. find its symmetry group. For this we reuse an implementation originally described in [11]. The elements of every symmetry group are isometries, and can be divided into translations, rotations, reflections and glide-reflections. It is easiest to look for these features in this order.




Figure 12: Example of the hyperbolization of an ornament

In order to find the grid of translations which map the ornament onto itself, one has to measure the similarity between the original ornament and translated versions of it. A naive implementation would consider every possible integral translation vector, and for each vector calculate a measure by comparing all overlapping pixels, e.g. by summing up differences in color values, as outlined in Equation 1. For a square image of edge length n this would yield an algorithm with $O(n^4)$ time complexity.

$$a(x, y) = \sum_j \sum_k |p(j, k) - p(j + x, k + y)| \quad (1)$$

A better approach is the calculation of the autocorrelation as a sum of products, not differences:

$$a(x, y) = \sum_j \sum_k p(j, k) \cdot p(j + x, k + y) \quad (2)$$

Using this definition of autocorrelation, the values for all possible translation vectors can be calculated in $O(n^2 \log n)$ by using fast Fourier transforms:

1. $P = \text{DFT } p$
2. $A(x, y) = \overline{P(x, y)} \cdot P(x, y)$
3. $a = \text{iDFT } A$

As Fourier transforms conceptually operate on repetitive data, the images have to be padded with zeros to at least twice their size in order to avoid any wraparound effects.

Once the autocorrelation has been calculated, the result can be searched for peaks. We employ an area of dominance approach to peak detection, which weights peaks by the area around them devoid of larger peaks. A similar approach to symmetry detection has also been presented in [17]. The heuristics used to extract basis vectors from this fuzzy

grid of peaks are still under development. Once two basis vectors spanning the grid of all translations have been found, they can be used to extract a single representative of a translative cell of the ornament.

In the next step, this cell can be tested for rotations and (glide-)reflections. We again employ correlations calculated using fast Fourier transforms, now between an image and its rotated or reflected copy. As there are only very few rotations which can occur in any Euclidean ornament, namely by 60° (6-fold), 90° (4-fold), 120° (3-fold) and 180° (2-fold), and also the direction of any axis of reflection has to correspond to certain vectors in the grid of all translations, there are only few correlations to perform. This time, we don't pad the image but instead make use of the cyclic nature of Fourier transformations. Because of this, we are not really correlating a single tile and its rotated copy, but rather an infinitely tiled area and its rotated version. From the position of a peak, the location of the possible center(s) of rotation or axis of translation can be derived. Also the glide distance of a glide reflection can be determined and used to distinguish simple reflections from glide reflections.

In general, the value of the peaks corresponding to each detected feature gives an indication as to how likely that feature is actually present in the ornament. In practice, however, it pays off to let the user decide about these features, as only the human can easily distinguish unintended errors from slight symmetry breakings intended by the artist.

From the set of detected features together with the grid of translations the concrete symmetry group of the ornament can be deduced. With this information, one can extract a single fundamental domain. For each pixel of the fundamental domain to be calculated, all the pixels in its orbit in the original image are accumulated and boiled down to a single color using averaging, median calculation, or a combination of these.

When used as a tile under the same symmetry group, this fundamental domain will yield an ornament that is similar to the input image but perfectly symmetric.

4.2 Reverse Pixel Lookup

Once the fundamental domain has been transformed, it can be used to fill the plane. The Poincaré disk model embeds the whole infinite hyperbolic plane into a finite Euclidean circle. To display this circle correctly, an infinite number of copies of the fundamental domain would be needed. As this is obviously impossible with finite computational means, one has to set some limits.

The easiest approach would be to simply create a finite number of copies, placing them one next to the other until some break condition is met. Doing so, one has to find a compromise between the completeness of the result and the time spent calculating it. Abort too early, and the image will have major holes along the boundary. Calculate too long, and most of the time will be spent transforming fundamental domains that are too small to be distinguished in any case. As we were dealing with interactive hyperbolic drawing applications before, and are planning to add a feature for drawing on the transformed ornament later on, we had quite severe real time constraints in mind for this tiling step, and thus this naive approach was deemed unsuitable.

The way to avoid this whole problem is to reverse the process. Don't copy tiles and see what pixels to color, but instead start with the pixels and look at the tiles they belong to. While there are infinitely many tiles in the Poincaré disk, every pixel of the disk belongs to exactly one such tile (neglecting pixels on boundaries), and every tile can be mapped onto a central tile with known color values by a finite number of semigroup generator transformations. For most of the area, two adjacent pixels will belong to the same tile, so that the transformation taking them into the central tile is the same as well.

This leads to the following algorithm, illustrated in Figure 13. Suppose you have a central fundamental domain with known content, illustrated as square pixels. You then copy colors from the central domain to the rest of the disk, moving from the center outwards. At any given point in time, some pixels are already colored (filled), while others are yet to be colored (outlined). Consider the next pixel to be colored (red circle, 1).

1. First apply the transformation (green arrow) which mapped its neighboring pixel (green circle, 2) into the central tile. The result (red circle, 3) is usually at least quite close to the central tile.
2. While the transformed position is still outside of any boundary line of the central tile (blue line), apply the generator associated with this boundary line (blue arrow). If the result (red circle, 4) still isn't inside the central tile, repeat this step until it finally is.
3. Round to integral pixel coordinates.
4. Copy pixel color.




Figure 13: Reverse pixel lookup in progress

5 A Picture Book

This section is dedicated to pictures created by the implementations of the methods described in the previous sections. Having well structured high quality input material turned out to be extremely important during the process of implementation. Most of the benchmarks came from the beautiful book “The Grammar of Ornament” published by Owen Jones in 1910. We now will present some of these originals together with interesting hyperbolizations.

5.1 Jones’ Ornaments

In Jones’ book there are various samplers of classical ornaments from very different cultures. The book was one of the first books produced with multiple color printing plates.

The novelty of the production process caused some nice irregularities and imperfections that were a challenge for our picture recognition algorithm. The following pages illustrate some of the most beautiful patterns of the collection together with several interesting hyperbolizations.

The first collection of images, Figure 14, shows in a sense *variations on a topic*. Each of the following eight hyperbolic pictures is generated from the same input material (upper left picture) and rendered with different indices for the rotation centers.




Figure 14: A sampler of hyperbolizations of a $p4$ ornament. Orders of rotation centers given in brackets.

The following pictures illustrate various other wallpaper groups with a small selection of hyperbolizations. Unfortunately it is not possible to represent the infinite variety of possible patterns within the finite space of this article. However we hope the reader will get an impression of the richness of the structures.

The first example Figure 15 shows an ornament based on a simple triangular reflection group (the $p6m$). Two straight forward hyperbolizations are shown.




Figure 15: Hyperbolizations of a $p6m$ ornament

Figure 16 has a $p2$ ornament as basis. The group $p2$ has four independent centers of two fold symmetry. Each of these centers can be altered independently. The pictures show the effect for two possible choices of a single rotation center. Notice that the $p2$ has a wide variety of possible appearances due to the freedom in the choice of geometric parameters.



Figure 16: Hyperbolizations of a $p2m$ ornament

Finally, in Figure 17 we present a hyperbolization of an even more subtle group: a pmg . Here in the original reflections, glides and 2-fold rotations are present. The group pmg has $p2$ as a subgroup. Thus the image can also be considered as a $p2$ with a particularly symmetric fundamental region. We used this fact to obtain the hyperbolizations by using our algorithm for $p2$.



Figure 17: Hyperbolizations of a pmg ornament


5.2 Escher's Lizards

It was a particular challenge and also particular fun for us to apply our algorithms to the ornamental work of M. C. Escher and by this creating new circle limit pictures similar to his. Here we will illustrate this process for a particularly interesting example: Escher's famous Lizard tessellation (symmetry drawing number 25). The original source for our hyperbolizations was Figure 18. In this case a particular problem arose. Escher's original drawing chose three different colors for the Lizards. If the colors are considered there is no rotational symmetry anymore. However, the shapes of the lizards admit a 3-fold rotation symmetry. We created a hyperbolization in a three step process. First we created a color free copy of the original Lizard picture, in which only the outlines of the Lizards were present (actually we created these outlines by hand using Escher's drawing as a blueprint). The picture of the outlines (which now admits a $p3$ symmetry) was fed to the hyperbolization algorithm. There we altered one of the 3-fold rotation centers to a 4-fold center and generated the hyperbolic image. Figure 19(a) illustrates the fundamental region of this ornament. The structure generated this way does no longer admit a proper three-coloring in a way such that no lizards of same color meet in a proper boundary segment. However, the picture is nicely four-colorable in a symmetric way.




Figure 18: M. C. Escher, symmetry drawing number 25, Ink, pencil, watercolor, 1939.

We then started to four-color the a fundamental region of this picture under the proposed coloring, as illustrated in Figure 19(b). This fundamental region was then used as an input for the reverse pixel lookup algorithm which finally produced a (four-colored) hyperbolization of the lizard picture. The result of this process is shown in Figure 20.



(a) Monochrome fundamental domain and its tessellation



(b) Colored fundamental domain

Figure 19: Intermediate steps when producing a hyperbolization of Escher's Lizards




Figure 20: Lizards reaching for the circle limit

6 What's Next

6.1 Subgroups

Many artistic ornaments almost belong to a group of higher symmetry, except for some details. E.g. the p4 ornament hyperbolized in Figure 14 would be a p4m except for the twist at the corners of the green almost-square. The shapes of Escher's Lizards from Figure 18 belong to p3, and only the colors reduce this to p1. It seems to be interesting to have large parts of an ornament stable under some highly symmetric group, but to have smaller details break this symmetry, resulting in some subgroup of the former group as the symmetry group of the whole ornament.

So far, we haven't dealt with such issues in depth. Our hyperbolization of the Lizards employs some code written for this specific purpose, which represents the specific subgroup used when colors are added. It would be nice to find and implement some general way to hyperbolize an ornament in a highly symmetric group and later on switch to a subgroup and intentionally break the symmetry. To transport artistic content and intention, the breaking of the symmetry should be related to what was done in the original image, but needn't be the same. So, at places where Escher broke his rotational symmetries using three colors, the hyperbolic ornament also breaks symmetries, although using four colors. The proper concepts to formalize symmetry breakings that don't interfere with the general hyperbolization still have to be found.

A first step in this direction would be an implementation in our interactive hyperbolic drawing application. We already used this application to develop many of the concepts discussed in this paper. One could allow users to select a subgroup, giving them a low-symmetry brush to modify an existing high-symmetry ornament, either hyperbolized or completely hand-drawn. Experiments of this kind would likely help understanding of the matter.

6.2 Spherical Geometry

Most of the concepts used for hyperbolization will work as well for spherical geometry. Especially the transformation of artistic content using discrete conformal maps and the reverse pixel lookup to render the resulting image should be valuable tools for a possible sphericalization toolkit. It would be interesting to implement such an application as well.

6.3 Generic Ornament Transformations

Of course one isn't necessarily forced to change geometry to conformally map an ornament. Even within the Euclidean plane there are symmetry groups with a common topology of their orbifolds, differing only in the orders of their rotation centers. Section 2.3 lists some such families. It would be interesting to allow conformal transformations between different Euclidean ornaments as well.

In theory, one could also start with a hyperbolic (or spherical) ornament and transform it to any other symmetry group. However, as the number of artistically interesting non-

Euclidean ornaments out there is rather low, such an application might have little practical value.

6.4 P1

In this article we completely excluded the groups $p1$, pm , pg and cm . They do not have any rotation centers. However it might be still desirable to translate their artistic content to a hyperbolic setup. In order to do so one must artificially introduce new centers of rotation. It might be an interesting study on the borderline of art and geometry to investigate what kind of places are suitable for such rotation centers. For instance if one has a wallpaper with a $p1$ flower pattern it might be undesirable to introduce a rotation center within some leaf or petal. A distinction between foreground and background will be important and except for rare cases rotation centers may occur preferably in the background.

6.5 Morenaments

The collaboration of the two authors of this article in the field of ornaments started with a real time drawing application for Euclidean ornaments, which has become an easy to use program for users of all ages and backgrounds, thus especially suited to public exhibitions. We followed this up with another application to draw hyperbolic ornaments, again in real-time. This is still being developed, and usability by laymen is still lower than we would wish. In the meantime, we have added further projects, like that for automatic pattern recognition or the tools for hyperbolizations discussed in this paper.

It is our goal to unite all these projects into a single application. The name of this combined application will be “morenaments”, in response to which we have already renamed the Euclidean and hyperbolic drawing applications to “morenaments euc” and “morenaments hyp”. We will try to achieve a user interface as easy to use as that of morenaments euc, while giving access to all the interesting features resulting from non-Euclidean geometries and the combination of different geometries. We will also use this integration to redesign most parts of the code, in order to accommodate features we have in mind but couldn’t implement yet.

Acknowledgements

The authors wish to thank all those persons who have been helpful to this project in its various stages. In particular we are grateful to Sasha Bobenko, Folkmar Bornemann, Boris Springborn and Yuri Suris for helpful mathematical comments. We are also very grateful to Mark Veldhuysen and the Escher Foundation for granting us permission to use artistic material of M.C. Escher in unmodified and in modified versions. We also wish to thank our partners Ingrid and Melanie for stylistic and artistic advice.

All M.C. Escher works © 2009 The M.C. Escher Company – the Netherlands. All rights reserved. Used by permission. www.mcescher.com

References

- [1] W. BARLOW. *Über die geometrische Eigenschaften homogener starrer Strukturen*. Z. Kryst. Min. **23** (1894), 1–63.
- [2] A. I. BOBENKO, YU. SURIS. *Discrete Differential Geometry: Integrable Structure*. Graduate Studies in Mathematics, Vol. 98, AMS, 2008, xxiv + 404 p.
- [3] C. S. KAPLAN and D. H. SALESIN. *Escherization*. SIGGRAPH 2000, the 27th International Conference on Computer Graphics and Interactive Techniques. New Orleans, Louisiana, USA, 25-27 July 2000.
- [4] H. BURGIEL, J. CONWAY and C. GOODMAN-STRAUSS. *The Symmetry of Things*. A K Peters Ltd, 2008.
- [5] H. S. M. COXETER. *The non-Euclidean symmetry of Escher’s Picture ‘Circle Limit III’*. Leonardo, **12** (1979), 19–25, 32.
- [6] H. S. M. COXETER. *The Trigonometry of Escher’s Woodcut ‘Circle Limit III’*. The Mathematical Intelligencer, **18** no. 4 (1996) 42–46.
- [7] T. A. DRISCOLL and L. N. TREFETHEN. *Schwarz-Christoffel Mapping*. Cambridge Monographs on Applied and Computational Mathematics, (2002).
- [8] D. DUNHAM. *More “Circle Limit III” Patterns*. Bridges London (eds. Reza Sarhangi and John Sharp), 451–458, 2006.
- [9] E. S. FEDOROV. *The elements of the study of figures*. [Russian] Zapiski Imperatorskogo S. Peterburgskogo Mineralogicheskogo Obshchestva [Proc. S. Peterb. Mineral. Soc.] (2) 21 (1885), 1–289.
- [10] E. S. FEDOROV. *Symmetry of Crystals*. Translated from the 1949 Russian Edition. American Crystallographic Association, New York, 1971.
- [11] M. VON GAGERN. *Computergestütztes Zeichnen in den Symmetriegruppen der euklidischen Ebene*. Diplomarbeit, TU München, 2007.
- [12] M. J. GREENBERG. *Euclidean and Non-Euclidean Geometries: Development and History*, 3rd ed. W. H. Freeman and Company, 1993.
- [13] M. HARMER. *Note on the Schwarz Triangle Functions*. Bull. Austral. Math. Soc., **72**(3), 2005, 385–389.
- [14] N. F. M. HENRY and K. LONSDALE. *International Tables for X-Ray Crystallography*, vol. 1, Kynoch Press, Birmingham, England, 1952.
- [15] O. JONES. *The Grammar of Ornament*. Folio edition, Bernard Quaritch, 1910.
- [16] J. LEYS. *Hyperbolic Escher*. http://www.josleys.com/show_gallery.php?galid=325.
- [17] Y. LIU, R. T. COLLINS and Y. TSIN. *A computational model for periodic patterns*. IEEE transactions on pattern analysis and machine intelligence **26** (2004), Nr. 3, 354–371
- [18] A. M. SCHÖNFLIES. *Über Gruppen von Bewegungen*. Math. Ann. **28** (1886), 319–342, and **29** (1887), 50–80.
- [19] A. M. SCHÖNFLIES. *Über Gruppen von Transformationen des Raumes in sich*. Math. Ann. **34** (1889), 172–203.
- [20] B. SPRINGBORN, P. SCHRÖDER and U. PINKALL. *Conformal equivalence of triangle meshes*. ACM SIGGRAPH, 2008.