

# Rendering 3D Tessellations with Conformal Curvilinear Perspective\*

Charles Gunn†

TU Berlin, Institute for Mathematics 8-3

10623 Berlin Germany

gunn@math.tu-berlin.de

## Abstract

Traditional perspective painting projects the world from the eye-point onto a single rectangular screen. In contrast, this article works with spherical images, created by projecting the world onto a closed sphere surrounding the eye-point, the *viewable sphere*. We focus on the use of stereographic projection of the viewable sphere to produce flat images which are conformal and include arbitrarily large fractions of the viewable sphere. We describe a practical, real-time implementation of this projection embedded in a general-purpose visualisation system. We then apply the method to provide a new way of visualising tessellations of three-dimensional euclidean (and non-euclidean) space, and give a preliminary discussion of the results.

**Motivation.** This work was inspired by the work of the American artist Dick Termes [10]. His desire to “paint the total picture” led him to paint on spherical canvases (Figure 1). The image he paints is that which an observer located at the center of the sphere would see. We call such an image a *viewable sphere*. Simultaneous with this artistic development, the last few years have seen an explosive growth in the use of such viewable spheres in photography, where they are known as *spherical panoramas*. [9] describes a “throwable ball camera”, 36 cameras packed in a spherical array, which generates such spherical panoramas with a single exposure. A natural question is “What sorts of flat images can be produced from viewable spheres?” This question is the impulse behind the present research.

**Looking out from the viewable sphere.** If you could position yourself at the center of one of Termes’ spherical paintings, then you would see what the artist saw when he painted the sphere. If you then insert a rectangular screen before your eye, perpendicular to the line of sight and project the sphere onto this screen, you would obtain a perspectively correct image of the original scene. In fact, all such traditional perspective images can be obtained in this way from a single viewable sphere by changing the direction of view and the size and shape of the rectangular screen. However, as you move away from the center of the sphere, such projections no longer map straight lines to straight lines. The series of images in Figure 3, all rendered with  $120^\circ$  field of view, illustrates how this distortion increases as one moves away from the center of the sphere, eventually arriving at the surface of the sphere.


---

\*An extended version of the article appearing in the printed proceedings of Bridges 2013.


†Research supported by Matheon Research Center for Key Technologies



**Figure 1 :** Two views of the spherical painting Cubical Universe by Dick Termes.



**Figure 2 :** Stereographic projection of sphere (only half shown) onto plane.




**Figure 3 :** *The inside view of the viewable sphere at distances 0, 0.25, 0.5, 0.75, and 1.0 from center.*


**Stereographic Projection.** The projection obtained when the eye or camera reaches the surface of the sphere itself, is a famous mathematical object known as *stereographic projection*, a conformal (angle-preserving) map of the sphere minus one point (the center of projection) onto the plane. Figure 2 illustrates the process of stereographic projection of three mutually orthogonal great circles on the sphere onto the ground plane. One half the sphere has been removed to aid the understanding. The projection of a point is found by joining the point with the north pole of the sphere with a line. Where this line cuts the plane is the projection point. If we widen our definition of *circle* to include lines (as circles of “infinite radius”), then stereographic projection is not just conformal, it also maps circles onto circles. We need the widened definition, since the image of a circle through the projection point is a line.

**Previous work.** [2] was a ground-breaking work which introduced the term “curvilinear perspective” and discusses practical and aesthetic issues involved in flattening spherical images. Flocon was inspired to undertake this work by his encounter with the work of M. C. Escher and the two engaged in a long-term correspondence. A more recent survey on the same theme is [3], which focuses on strategies for processing panoramic photographic images. Both of these sources present stereographic projection as a preferred solution but present interesting alternatives which are worth considering. The open software project Hugin ([hugin.sourceforge.org](http://hugin.sourceforge.org)) provides an environment for stitching together such spherical panoramas from individual photos and provides a variety of flattening projections. What is new in the present treatment is the description of a practical, real-time implementation for synthetic image generation, and the focus on the resulting visualisation of 3D tessellations in euclidean and non-euclidean spaces.

**Six-point perspective.** Termes describes his spherical paintings as “six-point perspective”, since all six cardinal directions – front, back, left, right, up and down – appear as vanishing points on the viewable sphere. In the interests of brevity, we employ the same term to refer to the images obtained from the viewable sphere



**Figure 4:** *Viewable sphere of a tessellation of euclidean space.*



**Figure 5:** *Six-point perspective: Stars mark vanishing points in 6 cardinal directions  $\pm x, \pm y, \pm z$ .*


via stereographic projection in this way. See Figure 5. The six star symbols represent the vanishing points. To be precise, the star symbol in the upper middle (lower middle) of the image represent the up (down) vanishing point. The left most star symbol is paired with the third from the left, as vanishing points at the end of a horizontal direction, say  $x$ ; while the rightmost and the third from the right star symbol are the vanishing points of the remaining pair of cardinal directions. See also Figure 6, which labels the six cube faces as  $(FT, BK, LF, RT, UP, DN)$ .

**From viewable sphere to six-point perspective.** Assume one has a viewable sphere represented in the computer as a sphere with the projection of the surrounding scene “painted” on it (typically in the form of *texture maps*). To obtain six-point perspective, one renders this sphere traditionally, using an on-axis camera positioned on its surface, pointed at the center of the sphere, yielding a stereographic projection. See Figure 7 which shows this rendering step using cameras with field of view of  $60^\circ$  and of  $120^\circ$ . The six-point perspective images shown in this article have been rendered with field of views of around  $120^\circ$ . Rotating the sphere while leaving the camera fixed produces a family of images in six-point perspective, typically with interactive frame rates.

**Viewable spheres generated by photography.** We illustrate this process of image generation using use a photographic spherical panorama provided by Jonas Pfeil of the Gendarmemarkt in Berlin, taken by one of his throwable cameras [9]. This data was provided as a cube map (see above), but other file formats such as the Android Photo Sphere (which uses a single polar coordinate image) can also be used. Note the typical “tiny planet” shape of the middle image. This is created by placing the projection point on the zenith of the viewable sphere (the point directly above the earth’s surface). Placing the projection point on the nadir point (beneath the earth) creates the inverse effect, seen on the right image.


**Generating viewable spheres using synthetic imaging.** In the absence of paint or photographs, one can generate viewable spheres using standard visualization systems. Assume one has some virtual scene represented in the computer. To create the viewable sphere, one positions the camera at the origin and generates a series of ordinary perspective images of the scene that can then be “sewn together” to create a viewable sphere. This is most simply achieved by rendering the scene six times onto the six faces of a cube centered on the eye point. That is, use an on-axis camera with a field of view of  $90^\circ$  and square viewport. These resulting images can then be texture-mapped onto the unit sphere. A set of six square images which go together to cover the viewable sphere is called a *cube map* and is used in computer graphics for a variety of effects (reflection maps, skyboxes, etc.) This method described here can be applied to any cube map. Naturally, one can also apply this technique to a ready-made cube map without generating the six images from scratch. See Figure 6, left.

**Using tessellations as an image source.** For this article we have chosen to focus on tessellations of three dimensional space as the source of synthetic imagery. We have found that tessellations and six-point perspective are good partners. Tessellations provide ideal examples to illustrate six-point perspective since



**Figure 6:** Viewable sphere of a tessellation of euclidean 3-space (left), standard perspective image (middle) and six-point perspective image, all rendered with cube face labels.





**Figure 7 :** *Stereographic projection of the viewable sphere using on-axis camera with field of view 60° (left) and 120° (right).*

the regularity of the tessellation helps the observer to understand the technique. In the other direction, the pictures accompanying the article demonstrate that six-point perspective provides a way of seeing the “whole tessellation” which ordinary perspective rendering cannot provide.

**Description of the tessellations.** Most of the tessellations in this article come from the 10 Euclidean manifolds (also known as platycosms) [1]. To understand fully the connection between manifolds and tessellations goes beyond the scope of this article. For mathematical details of the connection between the manifold and its group see [11]. For our purposes it suffices to know that every euclidean (non-euclidean) manifold is uniquely characterized by a discrete group of motions of euclidean (non-euclidean) space, and closed paths in the manifold correspond to group elements. Applying this to the optics of the manifold, the tessellation represents what an *insider* living in the manifold sees. Five of the 10 platycosms can be characterized in the following way: the group is generated by unit translations in the  $x - y$  plane, plus a screw motion in the  $z$ -direction that translates by 1 and rotates by  $\frac{\pi}{n}$  for  $n \in \{1, 2, 3, 4, 6\}$ . The other five are not so simple; four of them contain glide reflections. One way to characterize these 10 groups is as follows: the platycosm groups are those crystallographic groups whose symmetries have no fixed points.

We also include some non-euclidean tessellations. For this we take advantage of the fact that both hyperbolic and spherical space can be represented by projective models which map naturally onto modern visualization systems. The *120-cell* is a tessellation of the 3-dimensional sphere by regular pentagon dodec-



**Figure 8 :** *Outside view and two six-point perspective views of a viewable sphere of Gendarmemarkt, Berlin (image data courtesy Jonas Pfeil).*

ahedra with dihedral angles of  $\frac{2\pi}{3}$ . It has 120 group elements, hence the name. These copies are arranged in closed chains of 10; each pentagon dodecahedron belongs to 6 such chains (corresponding to its 6 pairs of opposite pentagonal faces). It is possible to choose 12 of these chains which do not intersect and which together fit together to make the 120-cell. It is analogous to the pentagon dodecahedron, which is a tessellation of the 2-sphere by 12 regular pentagons with angles of  $\frac{2\pi}{5}$ . As a partner space we present a hyperbolic tessellation whose tile is also a regular pentagon dodecahedron, but whose dihedral angles are right angles. These pentagon dodecahedra fit together 8 around a vertex, just like euclidean cubes, but the faces are pentagons, not squares.

For background on the visualization techniques used here for discrete groups see [4]; for the current state of metric-neutral visualization see [5].

**Visualization technique.** We create a polyhedron which serves as a fundamental domain, or tile, for the group, that is, it tessellates without overlapping. Of course if we tessellate with it naively we produce an opaque mass of tiles which fill space, preventing the viewer from seeing *into* the tessellation. To avoid this, we display the tile in two forms: first we thicken the edges into beams of variable thickness, and apply a woven texture map to this “skeleton”. Secondly, we include a scaled-down version of the faceted polyhedron. We often replace the latter with an asymmetric, texture-mapped alphabetic letter such as ‘d’, which due to its asymmetry tends to differentiate the types of symmetries better than the fundamental domain. Finally we create a list of group elements large enough (~5000) so that the tessellation appears infinite, shading off into the distance (enhanced by using fog effects in the rendering).

**Software.** The images were created with a plugin module built into the 3D, Java-based, metric-neutral scene graph package jReality [7]. It has been optimized so that the whole process is performed on the GPU. This allows multiple frames per second with a texture resolution  $6 \times 1024 \times 1024$ . For a fixed viewable sphere the frame rate is much higher. The reader can play with *SimpleManiview* [6], the program used to generate the figures in this article; it is available as a Java webstart application at the URL given in the citation.

**Discussion of figures.** As a general rule, the circles in the six-point perspective images correspond to straight lines in the original scene; a complete circle bounds the image of a plane in the original scene (see Figure 9, right). Figure 6 is based on the simplest tessellation of space,  $c_1$ , generated by 3 perpendicular translations. The corresponding manifold is the 3-dimensional torus. The three images show the viewable sphere rendered with a camera outside, inside, and on the surface (the latter is the six-point perspective image). In Figure 5 the same tessellation is featured but in black and white. Note that the two large light circles in the image (which intersect in two of the star symbols and contain the rest of the star symbols) correspond to the horizon line of the plane of the front left, resp. front right, face of the central tile of the




Figure 9 : Six-point perspective views of the euclidean platycosms  $c_6$  and  $+a_2$




tessellation. Each circle consists of vanishing points of sets of parallel lines in the corresponding plane.


Figure 9 shows six-point perspective views of two further platycosms. On the left  $c6$  is shown. As indicated above, it has a screw motion of order 6. The fundamental domain is a hexagonal prism which has been translated off-center. On the left side one sees a sequence of these prisms (showing a yellow face), spiraling around the axis of the screw motion. On the right, one sees the opposite face of the prism (red) as they spiral around the **other** end of the same axis. The figure on the right features the platycosm  $+a2$  whose fundamental domain is a cube, but whose symmetries includes glide reflections. Note again the distinctive dark circle, the horizon line of the plane containing the large square in the middle of the image. Figure 10 shows two further platycosms, both of which have a rhombic dodecahedron as fundamental domain. The group on the left is  $c22$ , which is generated by three screw motions whose axes lie on the face of a cube and are parallel to the three coordinate axes. The group on the right,  $-a1$ , also includes glide reflections. Note in both figures the vanishing points determined by the faces of the central rhombic dodecahedron, since the tessellation “stacks” infinite “towers” of cells over each face.

Figure 11 shows the standard three views of the viewable sphere for the 120-cell. The six-point perspective image (right) reveals much more of the intricate structure of how the 10-cell chains interweave than the standard perspective image. Finally, Figure 12 shows the hyperbolic tessellation by pentagon dodecahedra mentioned above. In the six-point perspective image one can see 11 of the 12 faces of the dodecahedron surrounding the camera. In both middle and right images one sees the characteristic feature of hyperbolic space that hyperbolic planes appear as disks in hyperbolic space, subtending a finite solid angle (relax your eye and look for circular forms). This is reminiscent of the effect noted in the previous paragraph, that the horizon line of a euclidean plane (sometimes called its *line at infinity*) can be mapped under stereographic projection to a finite circle in the image. The difference is that this happens in hyperbolic space also in standard perspective rendering. For further images, see the image gallery [8].


**Conclusion.** We have demonstrated a simple and effective implementation of conformal curvilinear perspective, which can be built into existing visualization systems. Although presented here primarily with synthetic imagery, we have also shown how to apply it to photographic imagery. We have applied this method to visualize a variety of tessellations of euclidean and non-euclidean 3-space, resulting in images that provide new insights into patterns and aspects not accessible via traditional perspective. This experience, along with the growing popularity of panoramic photography, leads us to suggest that six-point perspective may be a perspective projection particularly suited to the current *Zeitgeist*, since it provides a global, holistic view of (almost) the whole world rather than the limited, isolated rectangle provided by standard perspective



**Figure 10 :** Six-point perspective views of the euclidean platycosms  $c22$  and  $-a1$



**Figure 11 :** *Outside, inside, and “on” views of the viewable sphere of the 120-cell tessellation.*



**Figure 12 :** *Outside, inside, and “on” views of the viewable sphere of a hyperbolic tessellation.*

rendering.

**Acknowledgements.** The impulse for this work came from Dick Termes’ spherical paintings. The tessellations were inspired by Bill Thurston (1946-2012). Thanks to the referees for valuable comments.

## References

- [1] John Horton Conway and J. P. Rossetti. Describing the platycosms. *arXiv:math.DG/0311476*, 2006.
- [2] Albert Flocon. *Die Kurvenlineare Perspektive*. Medusa, Berlin, 1983. Originally published in French 1959.
- [3] Daniel M. German et al. Flattening the viewable sphere. In *Computational Aesthetics in Graphics. Visualization, and Imaging*, Banff, 2007.
- [4] Charles Gunn. Discrete groups and the visualization of three-dimensional manifolds. In *SIGGRAPH 1993 Proceedings*, pages 255–262. ACM SIGGRAPH, ACM, 1993.
- [5] Charles Gunn. Advances in Metric-neutral Visualization. In Vaclav Skala and Eckhard Hitzer, editors, *GraVisMa 2010*, pages 17–26, Brno, 2010. Eurographics.
- [6] Charles Gunn. Simplemaniview: A 3d manifold viewer, 2013. <http://page.math.tu-berlin.de/~gunn/webstart/jnlp/SimpleManiview.jnlp>.
- [7] jreality, 2006-2013. <http://www3.math.tu-berlin.de/jreality>.
- [8] <http://picasaweb.google.com/cgunn3/RenderItAll>.
- [9] Jonas Pfeil. A throwable ball camera, 2010. <http://jonaspfeil.de/ballcamera>.
- [10] Dick Termes. Painting the total picture. In *Bridges 2008 Proceedings*, pages 363–369, 2008.

- [11] William Thurston and Silvio Levy. *Three-Dimensional Geometry and Topology: Volume 1*. Princeton University Press, 1997.