MATHEMATICS MEETS PHOTOGRAPHY

Globe Workshop, by Lloyd Burchill

PART I

The Viewable Sphere

David Swart and Bruce Torrence

Right now, without moving from your seat or from where you are standing, look at your surroundings. Not just left and right but all the way around to whatever is directly behind you too. Look at the ceiling or the sky. Look at the floor below, or maybe it's a desk or a laptop below your nose. What you can see from your single point of view is a *viewable sphere*. Perhaps it helps to picture an imaginary sphere surrounding your head with imagery printed on it that matches your surroundings.

Of course, we need to refine this idea slightly in order to make it precise. A sphere has a single center, and you probably have two eyes. So stay very still, shut one eye, and imagine a sphere centered at the optical center of your open eye—the point where the light rays converge on their way to your retina. The radius of this imaginary sphere is not so important; let's just make it large enough so that your entire head is on the inside, and small enough so that it lies between you and every object within sight. Then for each point in the scene around you, even those behind you, the line segment from that point to the center of your open eye will intersect the sphere in a unique point. In this way, the scene around you can be painted, or at least projected conceptually, to create a well-defined viewable sphere.

To shoot an all-around panorama, a photographer takes a series of overlapping photographs in every direction from the exact same point in space.

With the advent of digital photography and continuously improving photo-stitching software, panoramas that capture the entire viewable sphere have become more and more commonplace.

To shoot an all-around panorama, a photographer takes a series of overlapping photographs in every direction from the exact same point in space—the optical center of the lens. It helps to have a wide-angle lens and a specialized tripod for this purpose. The photographer then imports these photos into software that can stitch them together into a seamless panorama.

These viewable spheres can be explored interactively. You might be familiar with applications such as Google Street View or panorama viewers that give online virtual walk-throughs of new homes, hotel rooms, or cruise ships. Others have created actual physical spheres with the panorama printed on them. Dick Termes, for instance, paints viewable spheres in what he calls "six-point perspective" because the viewable sphere contains all six vanishing points (two for each of three orthogonal directions).

However, a flat image is in many ways more accessible and practical than these computer-dependent applications. If we can find a mathematical mapping, or *projection*, that can map points from a sphere to a plane, then these viewable sphere panoramas can be printed and shared, hung on the wall, shown on a flat computer display, or included in a book or in *Math Horizons*. We need a function that takes points on the viewable sphere as input and outputs locations in the two-dimensional plane.

Luckily, cartographers and astronomers have been addressing this problem for millennia. What keeps cartography both interesting and difficult is a theorem by Riemann, which states that any mapping from a sphere to a plane will introduce some sort of distortion. Cartographers are naturally concerned with which sorts of distortions a projection introduces and whether a projection is suitable for the map-user's purpose. For instance, a specific projection might fail to preserve area, distances, angles, orientation, or any combination of these.

To get an idea of how projections are created or chosen for a purpose, consider an example. The *Mercator projection* is well suited for navigation because a straight line on a Mercator map corresponds to a path a ship would take while keeping its compass bearing constant. However, the Mercator map is *not* ideal if equal area is a concern because of the large changes in scale, especially near the poles.

Let's consider the *equirectangular projection* as a candidate to map viewable spheres. Even by cartography standards, it is old, attributed to Marinus of Tyre, circa 100 A.D.

Sébastian Pérez-Duarte, www.flickr.com/photos/sbprzd

Qualitatively speaking, conformal maps ensure that the imagery does not get sheared or skewed.

In this projection, the *meridians* (lines of longitude) are set as equally spaced vertical lines, and the *parallels* (lines of latitude) are equally spaced horizontal lines. For an equirectangular projection, you map a point from a sphere to a plane by simply renaming longitude and latitude to x and y, respectively. The left and right edges correspond to a single meridian line, while the top and bottom edges correspond to the North Pole and South Pole, respectively.

The format yields an image with a 2:1 aspect ratio because it takes 360 degrees to go all the way around the world longitudinally, but only 180 degrees to go from the North Pole to the South Pole. The equirectangular projection is so straightforward and useful that it has become the de facto standard native format for storing digital versions of viewable spheres.

A word of warning: In addition to cartography jargon such as meridians, parallels, equators, and poles, our discussion may include panorama terminology such as *zenith* (the point on the viewable sphere directly above the observer), *nadir* (the point below), and the *horizon* line.

Looking at the equirectangular projection above, we see that everything is oriented well (things that are pointed upward in the visible sphere are pointing upward in the equirectangular projection). However, as we can see, the equirectangular projection is not very well suited for viewing the parts of the panoramas near the zenith or the nadir. There is a rather nasty horizontal stretching at the top and the bottom of the projection.

Granted, these parts often comprise uninteresting flooring, grass, sky, or ceiling. But the features that *are* there are unrecognizable and can detract from the panorama's aesthetics. We want to look for projections that produce panoramas where no features are skewed or stretched. Do such projections exist? Under a reasonable interpretation of the phrase "no features are skewed or stretched," the answer is a resounding yes!

Conformal Mappings

When looking at different attributes of projections, one property seems to be better suited for photographic content than others: *conformality*. Mathematically speaking, conformal mappings are mappings that preserve angles at a local level. For instance, if two curves meet at, say, a 45-degree angle on the viewable sphere, then their images in the plane under a conformal mapping will also meet at 45 degrees.

To be conformal, a mapping must preserve every angle on the viewable sphere.

Qualitatively speaking, conformal maps ensure that the imagery does not get sheared or skewed; there is no squishing and stretching of the sort that we see near the top and bottom of an equirectangular projection. (While delightfully simple, equirectangular projections are *not* conformal.) As a consequence of preserving angles, it can be shown that any squishing or stretching that does happen under a conformal mapping occurs in all directions uniformly. This ensures that small features retain their shapes, even though we can expect some variations in scale and orientation. In other words, small-scale details will keep their overall shape and appearance, but large-scale shapes might grow, shrink, bend, or become distorted in some manner. This change of scale is not as problematic for photographic content as it is for maps perhaps because your eye tends to accept larger and smaller details as closer and farther away, respectively.

The conformal *stereographic projection* is about as old as the equirectangular projection. It is attributed to Ptolemy in the second century A.D. It is one of the most popular projections for viewable spheres. If we imagine the viewable sphere itself as a translucent ball resting on a white floor, we could place a light bulb at the zenith of the sphere and then look at the floor. The imagery on the floor is the stereographic projection of the sphere. (See figure 1). This mathematical function lives up to its name as an actual projection.

The stereographic projection of a viewable sphere is striking. Everything below the horizon is transformed to the interior of a circular region reminiscent of the little planet from Saint-Exupéry's *Little Prince*. So it is no surprise that these projections have been dubbed "Little Planets."

Sébastian Pérez-Duarte, www.flickr.com/photos/sbprzd

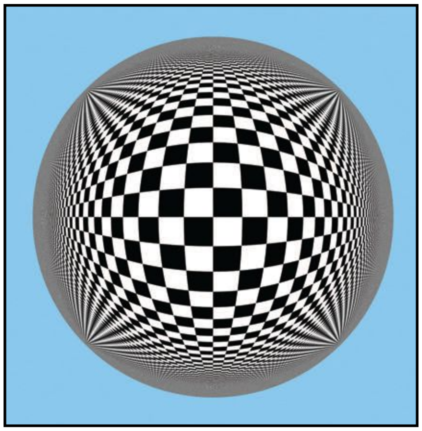

Figure 2. *Stereographic projection of a viewable sphere centered above an infinite checkerboard plane.*

In a stereographic projection, the nadir of the panorama is moved to the center of this little planet sphere. Because the nadir is so prominent, it is not uncommon to see panoramas that are taken from a point directly above a decorative floor element.

The horizon, which forms the equator of the viewable sphere, gets projected onto a circle. Everything below the horizon lands inside the circle, and everything above the horizon lands outside it. The photographic details at the horizon are such that the horizontal plane of the ground underfoot is parallel to the viewer's line of sight. So when this horizon is mapped to a circle, it gives the appearance of being the edge of a sphere viewed head on, as you can see in figure 2.

Figure 1. *Stereographic projection transforms the spherical image to a flat one.*

Lines to Circles

Stereographic projection can be better understood by seeing what it does to something as simple as a line in the original scene. What happens to the edge of a sidewalk or building, or to the edges of the magazine that you are now reading?

Recall that we first project the three-dimensional scene onto the viewable sphere. We then apply stereographic projection to the spherical image. We need to follow a line in the original scene through these two stages. For the first stage, start with any line in the scene that does not contain the center O of the viewable sphere. It will be projected to a *great circle* on the viewable sphere, as shown in figure 3. To see why, simply note that the line together with the center O determine a plane. Since this plane passes through the center of the viewable sphere, it intersects the sphere in a great circle.

For the second stage, we are left with the task of determining what happens to a great circle under stereographic projection. We will see that there are two distinct cases: those great circles on the sphere that pass through the zenith and those that do not.

The first case is simple. Any great circle containing the zenith must also contain the nadir. Such circles are meridians; as such, they correspond to *lines* under stereographic projection. Every such line passes through the nadir at the center of the final image, and we could reasonably call them *radial* lines. Since *vertical* lines in the original three-dimensional scene are portions of meridians on the viewable sphere, we conclude that vertical lines in the three-dimensional scene map to radial lines in the stereographic projection. More

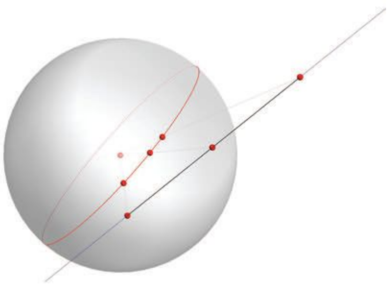

Figure 3. *A line in the scene becomes a great circle on the viewable sphere.*

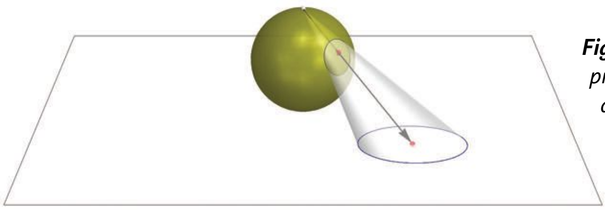

Figure 4. *Stereographic projection maps circles that do not contain the zenith to circles.*

generally, any line meeting the horizon at a right angle corresponds to a meridian on the viewable sphere, and thus maps to a radial line in the stereographic projection. Can you see the radial lines in figure 2?

The second case is slightly less simple. It can be shown that any circle (great or otherwise) on the surface of the sphere, that does not pass through the zenith, will map under stereographic projection to another *circle*. See figure 4, and consult the Needham reference in the Further Reading section for details. This implies that lines in the original scene that are not meridians on the viewable sphere will map stereographically to *circles*. Lines that intersect the horizon at nearly a right angle (great circles that are nearly meridians on the viewable sphere) will map stereographically to large circles. Lines that are closer to horizontal will map to smaller circles. The smallest circle in the stereographic projection that corresponds to a line in the original scene is the horizon circle; no great circle on the sphere projects stereographically to something smaller.

So we have answered the question of what happens to linear features in the three-dimensional scene. Every straight line in the original scene becomes, under stereographic projection, either a portion of a radial line or some portion of a circle, according to whether or not the original line corresponds to a meridian on the viewable sphere. Moreover, it is easy to see that any sphere projects in a circle on the viewable sphere and maps to a circle via stereographic projection.

Now close one eye again, and look for lines in the scene around you. Which of them correspond to meridians on the viewable sphere, and so become radial lines in the stereographic projection, and which are skew to the horizon, and so become circles instead? Can you imagine what the little planet version of your current surroundings would look like? See The Zip-Line on page 30 for a stereographic visualization challenge.

Stay tuned for the second part of this article, to appear in the next issue. We will address a host of other conformal projections in a systematic way by making use of a seminal idea of Riemann's: we identify the standard stereographic image with the *complex plane*. By composing stereographic projection with a second conformal map from the complex plane to itself, myriad visual effects are possible. ■

Further Reading

Examples of panoramic imagery abound at Flickr.com; search for "stereographic" or "equirectangular."

Carlos Furuti maintains an excellent website cataloguing properties of various cartographic projections: visit *www.progonos.com/furuti/MapProj/CartIndex/cartIndex.html*.

The geometry of stereographic projection, and in particular the fact that circles map to circles, is explained beautifully in chapter three of Tristan Needham's book, *Visual Complex Analysis* (Oxford University Press, 1999).

If you want to try making your own panoramas, the (free, open source) Hugin stitching software is a good place to start. Smartphone apps such as Autostitch and Photosynth are gaining popularity and are very simple to use.

—David Swart and Bruce Torrence

David Swart is Advanced Product Developer at Christie Digital Systems. He has a master's in mathematics and specializes in software that warps and blends multiple projectors into seamless displays. **Email:** david.swart@christiedigital.com

Bruce Torrence is the Garnett Professor of Mathematics at Randolph-Macon College. He attempts to warp and bend the minds of his students.
Email: btorrenc@rmc.edu

http://dx.doi.org/10.4169/mathhorizons.19.1.14