

Flattening the Viewable Sphere

Artistic Submission

D. M. German¹, L. Burchill², A. Duret-Lutz³, S. Pérez-Duarte⁴, E. Pérez-Duarte⁵ and J. Sommers⁶

¹ dm@uvic.ca, ² lloyd@flamingpear.com, ³ Alexandre.Duret-Lutz@lip6.fr,
⁴ sebastien.perez-duarte@m4x.org, ⁵ emmanuel@perez-duarte.net, ⁶ josh@sommers.net

Abstract

The viewable sphere corresponds to the space that surrounds us. The evolution of photography and panoramic software and hardware has made it possible for anybody to capture the viewable sphere. It is now up to the artist to determine what can be done with this raw material. In this paper we explore the underdeveloped field of flat panoramas from an artistic point of view. We argue that its future lies in the exploration of conformal mappings, specialized software, and the interaction of its practitioners via the Internet.

1. Introduction

Look around you. Barring the effects of stereoscopy, we perceive our surroundings as a sphere centered somewhere inside our head: the viewable sphere. It is generally accepted that our eyes perceive the world in a similar way in which a pin-hole camera records a scene [Kub86]. This is equivalent to mapping a section of the viewable sphere (usually a region of field-of-view of at most 120°) into a flat surface using a perspective projection (also known as gnomonic, or rectilinear). The perspective projection has a crucial feature: it preserves straight lines as straight. Its main disadvantage is that, as the field-of-view increases, the regions at its edges are heavily distorted. An image of field-of-view of 180° would have infinite size.

Mapping a sphere (or a large portion of it) is a problem well-known to cartographers who have been interested in making flat representations of the world [Sny93]. They have envisioned dozens of different map projections that take the sphere and project into a plane [SV89]. The choice of projection is a trade-off: one has to accept some distortion in exchange for some properties that the projection exhibits. For example, the Mercator Projection, one of the most famous, has been very valuable for navigational purposes because it preserves angles at any point, even if distances were distorted. A sailor would know exactly in which direction to set sail, even if he or she would not know how far the destination was.

The first panoramas (representations of the viewable sphere or a large portion of it) were developed by the Egyp-

tians and the Greeks, who used the stereographic projection for star maps [Sny93]. During the Renaissance artists developed the theory of projection, and understood its limits. It was Robert Baker, in the late 1700's who is credited as the father of panoramas and who used a cylindrical room to display 360° views of London with little distortion.

Photography, by its own nature, has been restricted to the optical limitations of the camera obscura and lenses (rectilinear, fisheye, and anamorphic). To overcome such limitations cinematographers developed sophisticated systems, such as Cinérama, Circarama, and Circle Vision 360° that uses 9 or 11 cameras and a cylindrical screen to project an image of 360° horizontal field-of-view, but less than 90 degrees vertical [Mac57]. Photographers had to accept the limitations of their medium, and present wide-angle images with heavy distortions towards the edges; yet many used this feature as an artistic tool.

The computer has open a new era in the creation of panoramas. Affordable hardware and software exists to capture and create a $360^\circ \times 180^\circ$ panorama. These spherical panoramas are usually displayed as immersive panoramas (using software such as Apple Quicktime VR) and have found an important market in surveillance and real estate, by providing a "realistic" view of a space as if one was there. Yet some photographers are interested in displaying their images flat, in a way that they can be displayed without the need of artifacts that detract from the image, or limit its access.

This article is written by artists who are interested in this

challenge: how to represent the viewable sphere in a flat representation. In section 2 we introduce the notion of conformal mappings which we have found to be extremely useful. In section 3 we explore the use of the stereographic projection, the oldest of the map projections, but also, and until recently, the only conformal projection available in panoramic software. In section 4, we explore how Flickr—a Web site to post and view photographs—is becoming a gathering place and a laboratory where new ideas are tested, shared and learnt. Finally, section 5 explores the merge of the software developer, the mathematician, and the artist, and how these skills come together to create software that can unleash the imagination of a photographer.

2. On conformal mappings

Complex numbers. These simple words would usually make a few artists shudder, with memories of a long gone (and perhaps thankfully so) past. However it is with tools taken from this part of mathematics that graphic artists could fine a goldmine of possibilities for the transformation of images. Complex functions are a deeply profound theory with extremely powerful theorems, and, compared to what can be done, it is virtually unexplored. The development of fast computation can allow for an easy and affordable exploration and use of these functions. Before dipping into the sea of possibilities, we should present a quick introduction of what makes complex analysis so interesting for image transformations.

2.1. Complex numbers and complex functions

Everything starts with the quantity $z = x + iy$, where i is the imaginary quantity “square root of -1”. x and y are real numbers, called the real and imaginary part of z , respectively. The relationship between images and complex numbers is simple: the x and y can represent the coordinates of a point (x, y) in what is known as the “complex plane”. Any transformation of a complex number can thus be seen as a transformation of an image.

On this complex number we can define, just as for real numbers, functions, for example $f(z) = z^2 = x^2 - y^2 + i2xy$. The result of applying this function to an image would be to send the point (x, y) to the point $(x^2 - y^2, 2xy)$. All sorts of functions can be used: in fact, almost all usual functions have a natural extension to the complex plane. Sinus, cosinus, square root, exponential, logarithms, inverse trigonometric functions,... all have their counterpart in the complex world. In order to experiment with such functions we use Mathmap, a plug-in for the Gimp [Pro07]. Mathmap defines its own domain-specific programming language, which is then used to implement the desired remapping function (given the way the image is rendered one needs to program the inverse of such function, which is not always trivial). Mathmap provides a flexible and powerful environment for the exploration of functions to remap an image.

2.2. Derivatives and conformality

These function can be differentiated, just like real functions: $f'(z) = \lim_{\epsilon \rightarrow 0} \frac{f(z+\epsilon) - f(z)}{\epsilon}$. This derivative can only be defined if the real and the complex parts of f are related in some specific sense. This means that differentiation is in fact more constraining in the complex world than in the real world. Such a differentiable function defines on the complex plane “a conformal mapping”. The characteristics necessary for the function to be differentiable provide the fundamental property of conformal mappings: conformal mappings are shape-preserving, or angle-preserving, i.e., a transformed angle measures the same, i.e., local shapes are preserved [Neh82].

This is not to say that conformal mappings show no distortion: the scale of objects is distorted, but their shapes are preserved (and recognizable) regardless of where the object appears in the image. Sometimes, conformality fails in one point: the function is not differentiable there (example: $f(z) = 1/z$ in $z = 0$). At this point the preservation of angles is not observed, and the distortion is maximal.

A good example of a conformal mapping is the so-called “Escher Droste effect”, obtained by using the complex function $f(z) = z^{1+i\alpha}$ with well chosen values of α (see figure 1) [dSJ03]. This will transform the image into an infinite recursive spiral that shows nowhere the signs of having been distorted. Figure 2 shows an example of the polynomial transformation $f(z) = z^2 + 2z$.



Figure 1: Honey, I Escherized the Kids! Image created using the Droste effect. ©Sébastien Pérez-Duarte, used with permission.

2.3. The power of conformality

Incredible results are possible: Riemann stated in 1851 (and this was later shown to be true) that any reasonable simply-connected shape can be conformally transformed into a disk. Then Schwarz and Christoffel independently found an explicit but computationally intensive formula to transform the



Figure 2: Smarties: Heart Mapping. Applying the conformal mapping $f(z) = z^2 + 2z$; the plate was transformed from a circle to a heart, while preserving the shape of the smarties. ©Alexandre Duret-Lutz, used with permission.

unit circle s into any polygon w of n sides [Lee76]:

$$w = \int_0^z (1 - s^n)^{-2/n} ds$$

There are, however, known and simpler mappings between the disk and a square, a triangle, a rectangle, an ellipse using Dixon and Jacobian elliptic functions [Lee76]. With a fast computer and appropriate techniques, almost all shapes can be transformed into other shapes, keeping always those precious angles constant.

2.4. Conformality and the impossible flat sphere

In the accompanying exhibit to the Symposium *The Viewable Sphere* conformal mappings have been used to deal with the frustration on the impossibility of trivially flattening the sphere. How can we represent as truthfully as possible the globe? One logical answer is to use a projection that is shape-preserving, in the local scale. Here of course enter the conformal mappings between the sphere and the plane.

The first conformal projection known is the stereographic projection [Sny93]. This projection shows only one point where conformality fails: the opposite pole. The sphere minus one point is projected onto a plane of infinite size. The second conformal projection is well known: the Mercator projection is also angle-preserving. Distortions are larger close to the poles, where conformality fails. The sphere is projected into an infinite band.

It is also possible to conformally represent the whole sphere into an hemisphere, and this hemisphere can be projected as a circle by a stereographic projection (Lagrange projection). This circle can then be transformed into a square thanks to Schwarz's formula (Adams World in a Square). It is also possible to independently project each hemisphere into a disk and square this disk (Peirce quincuncial mapping and Guyou Doubly Periodic map—see figures 3 and 4).



Figure 3: Peirce Quincuncial Projection. ©Sébastien Pérez-Duarte, used with permission.




Figure 4: Tileable Guyou-Peirce projection. ©Sébastien Pérez-Duarte, used with permission.

In this very dynamic environment, new transformations are just waiting to be discovered.

3. On the Stereographic Projection

The Stereographic Projection is the oldest conformal mapping known. It projects the surface of a sphere on a plane that is tangent to the sphere's pole. To calculate the projection of a point P of the sphere, imagine a straight line between the opposite pole of the sphere (center of the projection) and P: the point where this line intersects the plane is the stereographic projection of P. The hemisphere tangent to the projection plane is thus mapped into a disc that is twice as big as the equator of the sphere, while the other hemisphere fills the rest of the plane.

When used to display $360^\circ \times 180^\circ$ panoramas, this projection yields pictures with a planetoid look such as in figure 5,

or, if you switch the two poles, that have a tunnel feeling as in figure 6. In any case the center of the projection is never seen, because it would be projected on the plane to the infinite (in all directions).



Figure 5: Notre-Dame de Paris, using the stereographic projection where the plane of projection is at the nadir, and the point of projection at the zenith. ©Alexandre Duret-Lutz, used with permission.



Figure 6: Anti-G Tunnel, using the stereographic projection where the plane of projection is at the zenith, and the point of projection at the nadir. ©Emmanuel Pérez-Duarte, used with permission.

Being conformal, the projection preserves local angles; but what makes the projection unique is that it will also preserve circles and vertical lines. Any circle on the original sphere (that doesn't pass through a pole) will be projected as a circle on the plane: this is a property few other projections share. The meridians of the sphere (circles that pass through the center of projection), that support the vertical lines of the scene photographed, will appear as straight lines coming

out of the origin of the projection plane. In photographs, it means footballs will stay round (unless they are American) and the vertical edges of buildings will remain straight.

The projection can be tweaked by moving the projection plane, or rotating the original sphere. Moving the projection plane in parallel, to change its distance to the source of the projection (i.e., so it either intersects the sphere at some constant latitude, or doesn't touch the sphere), will only affect the scale of the projection. This can be used to zoom in or zoom out the panorama to frame it properly.

When rotating the sphere, the center of the plane will no longer be located at a pole (the zenith, or the nadir). This means that vertical lines of the original scenery are no longer preserved. Off centering the projection a bit can give a jelly look to buildings, as in figure 7. It is also an option when both poles should be shown.




Figure 7: Planet Grain Mill, using an oblique stereographic projection. Many of the vertical lines appear curved. ©Josh Sommers, used with permission.

4. On Flickr

Art is, by its own nature, evolutionist. Artists have always been influenced by their environment and the artistic styles of others. Paris, for example, has an ample tradition incubating art. Monet, Renoir, Basille and Sisley were students of Gleyre, who was a student of Delaroche. Cézanne, Manet, Monet, Pissarro, Renoir, Sisley, Gauguin and Degas knew each other and their art. In the 1800's and early 1900's Montmartre, in Paris, became one of the most important artistic centers in the world. Matisse, Van Gogh, Renoir, Degas, Toulouse-Lautrec, and Picasso had studios there. For many artists Paris became the place-to-be.

In the past, communication was slow, hence, artists were more likely to move to the places where they could improve their art, and find potential buyers. The Internet has revolutionized the way we communicate, and the way that artists

interact. Artists started using mailing lists to communicate; as the Web developed, specialized sites were created for this purpose.

In the past we felt lucky if a few people were able to see our photos, and even then it was only because we would show a friend or family member a photo on a computer, or hanging on a wall. At first sight Flickr is a Web site to share your snapshots with your friends and relatives. But it is in its social interaction, and its ability to link strangers that its real power lies.

Flickr has permitted the creation of groups of photographers in particular niches of interest to develop. Among those niches, we think of course about the spherical panoramas and their multiple representations. Gathering around this wonderful new tool, artists around the world are able to share their art, and their techniques with others. The resulting artistic world is therefore one of greater diversity and originality—an artistic world in which everyone can find his own “little planet”.

To explore other’s photographs on Flickr is a unique experience. The vast collection of images on Flickr seems seemingly infinite and we soon started to stumble across images that blew our mind. Many of the images we found on Flickr have become a source of fascination, and more important, inspiration. We post our images regularly and from the work of others, we explore and learn new techniques at a very rapid pace. As time went on, our photo streams have attracted the attention of other Flickr members and our work is getting attention from all over the world.

Each of the Flickr groups are communities and that comes with all of the benefits of any other community, and yet, Flickr allows its members to move across communities, visiting, and in many cases joining and participating in a particular one.

Flickr is a place to share your artwork with peers from around the world, and see the work of others from around the world too. It is a place of learning and of teaching, a place to inspire and be inspired. For us, Flickr is perhaps the only way that our work could have been seen by so many people, from so many places in such a short amount of time.

This is the beauty of Flickr. It is so much more than a place to post photos. Flickr groups are the Internet’s Montmartres.

5. Flexify: Building software with an artistic vision

By *Lloyd Burchill, Flaming Pear Software*
<http://www.flamingpear.com>.

Computers can grasp the formal properties of pictures—shape, colour, perhaps 3D scene structure—but have next to no ability to understand their meaning. Flaming Pear’s challenge in building graphics software, then, is to emotionally transform pictures via formal manipulation.

One way software can do this is to steal techniques from traditional painting and photography. For example, a crisp photo can be made remote and moody by softening focus, vignetting edges, subduing color, reducing contrast, and adding grain. Less explored, and more interesting, are manipulations which only software can do and which let people see the world in new ways. To do this Flaming Pear pilfers methods from medical imaging, satellite photography, and scientific image processing and perverts them to aesthetic ends. Cartography in particular offers much to plunder for panoramic photos.

After the technical hurdles of building a spherical panorama comes the problem of how to present it acceptably to human eyes. Though there’s little tradition of fine-art spherical imagery to draw upon, there’s some luck: flattening a photo-sphere onto paper is the same problem as making a flat map of the world.

Flaming Pear’s panorama software Flexify, a plug-in for Adobe Photoshop built upon Adobe’s plug-in SDK, can warp equirectangular panoramas into traditional map “projections”. Typically these maps are designed to suit conditions which don’t much benefit a photograph: they avoid interrupting continents, or they simplify marine navigation by showing straight rhumb lines.

In response we’ve tried to invent several new projections which do suit photos. These new views aim either to present the whole scene clearly, or to draw out the vertiginous, dizzying, hyper-wide-angle quality of spherical panoramas.

Such panoramas are often beset by serious compositional problems: a sparsity of interesting features or lack of a clear focus for the viewer’s gaze. Ordinary pictures solve this with cropping, but to discard part of the picture misses the point of a complete spherical view. Ransacking some obscure, unpopular map forms from the late 19th and early 20th centuries allows the creation of unusual panoramas which improve the compositional problem. For example, the Lee Tetrahedric map [Lee76] applied to a photo produces the St Dunstan image shown in figure 8

These are conformal projections: they preserve local shapes but greatly distort scale. World maps made this way tend to be objectionably warped, but a photograph processed the same way is much more pleasing. The scale variation which is a flaw cartographically is an asset photographically since it allows the user to magnify interesting parts of the image and shrink the rest. So this type of projection is the one most favoured by panoramic photographers.

The abandonment of cartographic constraints in favor of aesthetic freedom has complicated Flexify’s underlying math. The Peirce Quincuncial and other polygonally-shaped projections involve the Schwarz-Christoffel transform (see section 2.3). Calculating the inverse transforms—from the output projection back to the disk, which in turn trivially maps to a sphere—is slow. A typical 4-megapixel image can



Figure 8: St Dunstan. Spherical panorama projected using the Lee Tetrahedric Map projection. ©Lloyd Burchill, used with permission.

take several minutes on a current 2 GHz consumer PC. This can be greatly accelerated by breaking up the output polygon into triangles whose barycentric coordinates are used as indices into a lookup table calculated offline. In practice, coarse tables of about 50x50 entries, linearly interpolated, reduce render times from minutes to seconds with no perceptible loss of image quality. High-quality resampling is achieved through the Feline algorithm [MRPaJ99] and sinc filtering [Tur90].

Despite these issues the software is quite simple to use. The user loads an image into Photoshop, starts Flexify, and through pop-up menus describes the input image supplied and can cycle through the different output projections with the help of a low resolution preview. Sliders tumble the image sphere, letting the user rapidly evaluate various options and discover the projection which flattens the photograph best. Figure 9 shows Flexify in action.

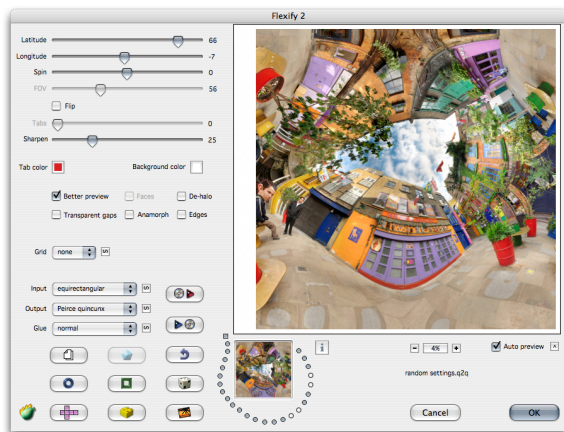


Figure 9: Screenshot of Flexify's graphical user interface

The software is a paintbrush, not a painting, and tools are not talent, so the creative onus lies now as ever with the human user. But it helps to have good equipment with strange new capabilities.

6. Conclusion

We feel like trailblazers exploring new artistic worlds. One day panoramic images created using conformal mappings might share wall space in galleries and museums along some masterpieces of our civilization. In the meantime we are digging into the treasure chests of geometry, geography, cartography, complex number theory and psychology (to name a few fields) trying to find tools that will lead us towards novel ways in which we can represent the viewable sphere.

We invite the viewer to visit the accompanying electronic exhibit for CaE 2007: *Flattening the Viewable Sphere* at <http://turingmachine.org/viewableSphere> where we showcase our images and those of other panorama artists.

References

- [dSJ03] DE SMIT B., JR. H. L.: The Mathematical Structure of Escher's Print Gallery. *Notices of the AMS* 50, 446–451 (2003).
- [Kub86] KUBOVY M.: *The Psychology of Perspective and Renaissance Art*. Cambridge University Press, 1986.
- [Lee76] LEE L.: Conformal projections based on elliptic functions. *Cartographica*, Monograph no. 16, supplement no. 1 to *Canadian Cartographer*, v. 13., 1976.
- [Mac57] MACGOWAN K.: The Wide Screen of Yesterday and Tomorrow. *The Quarterly of Film Radio and Television* 11, 3 (1957), 217–241.
- [MRPaJ99] MCCORMACK J., R. PERRY AND K. F., JOUPPI N.: Feline: fast elliptical lines for anisotropic texture mapping. In *Proceedings of ACM SIGGRAPH* (1999), pp. 243–250.
- [Neh82] NEHARI Z.: *Conformal Mapping*. Dover Publications, 1982.
- [Pro07] PROBST M.: The MathMap GIMP Plug-In. <http://www.complang.tuwien.ac.at/schani/mathmap/>, 2007.
- [Sny93] SNYDER J. P.: *Flattening the Earth, two thousand years of map projections*. University of Chicago Press, 1993.
- [SV89] SNYDER J. P., VOXLAND P. M.: *An Album of Map Projections, Professional Paper 1453*. U.S. Geological Survey, 1989.
- [Tur90] TURKOWSKI K.: Filters for common resampling tasks. 147–165.