

New methods to project panoramas for practical and aesthetic purposes

Daniel M. German¹, Pablo d'Angelo², Michael Gross and Bruno Postle

¹dmg@uvic.ca, Dept. of Computer Science, University of Victoria. ²pablo.dangelo@web.de

Abstract

Recent advances in digital photomontage have simplified the creation of extreme wide-angle views from a vantage point, including the recreation of the entire sphere (we will refer to these type of images as panoramas). In order to minimize the distortion from the point of view of the viewer, panoramas have been typically presented using curved displays (such as the original panoramas, by Barker, in 1787; or several cinematographic systems, such as Circle-Vision 360, still in use), and more recently with the help of the computer (such as the QuickTime VR format). Unfortunately requiring such systems restricts their use, and little research has been done in the representation of panoramas into a flat surface. In this paper we propose the use of several geographic map projections to project a panorama into a flat surface, both for realistic purposes (where the projection can be easily accepted as a faithful representation of the original image) and for artistic purposes (where the projection is used as an artistic tool intended for the creation of an innovative interpretation of the panorama). Finally we explore the use of inclinometers and map projections to automatically project an image from a wide-angle lens (rectilinear or fisheye) into a new image that is more aesthetically pleasant.

We believe the projections discussed in this paper will be useful to photographers, artists, and the designers of virtual reality environments, all of who might require the displaying of images with a wide field-of-view.

1. Introduction

The viewable sphere is the space around as we perceive it, centered somewhere inside our head. Our physical limitations do not allow us to capture the entire sphere at once. Our horizontal field of view (H-FOV) is close to 180°, although for practical uses it is smaller. Our vertical fields of view (V-FOV) is smaller than the H-FOV and it is restricted by our bone structure. In order to perceive the viewable sphere we need to rotate our heads.

Artists have always been interested in ways to represent part, or all the viewable sphere in a believable manner. Renaissance painters were the first to understand and use perspective. They also understood that there were limitations in the use of perspective that were not easy to overcome: as the FOV of a scene increases, its distortion towards the edges increases rapidly [Kub86]. In 1787, Robert Barker presented the first panoramas in semi-cylindrical rooms. The visitor had to stand in the middle to fully appreciate them. Recent advances in panoramic hardware and software have made it

easy to capture and recreate the complete viewable sphere (360°x180°FOV) or a large portion of it. But projecting this sphere into a flat surface is not trivial. This is one of the main reasons that photographic panoramas are displayed using special software (such as Quicktime QTVR) that permits the viewer to pan and zoom inside this sphere as if she was in the middle of it. Requiring a computer or special rooms (as in the case of Barker) has constrained the uses of panoramas. Flat panoramas are desirable because they do not require any special equipment or artifact to be displayed.

Projecting the sphere or part of it has been a problem faced by cartographers since the Greeks and Egyptians discovered that it was difficult to create flat star maps of the celestial sphere [Sny93]. Cartographic uses are very different from photographic use. Perhaps the most important difference is that, for most uses, panoramas do not need to be precise, as they do not need to be relied upon for navigation nor for land surveying. Panoramic software (such as panotools and Hugin—the authors of this paper are the main maintainers of both applications) has concentrated in the creation of the

viewable sphere, and has provided a limited number of projections, such as the equirectilinear, the central cylindrical (Barker probably used a central cylindrical projection for his paintings), and cubic rectilinear (the sphere is mapped into a cube, where each side is projected using a rectilinear projection), and the sinusoidal. Recently the stereographic projection was introduced and has been widely used for artistic purposes. As in cartography, the choice of a projection is a trade-off: some benefits are gained, in exchange for some distortion. A panorama-maker (like a map-maker) has to determine what are the features that are more important to preserve in a panorama, and choose the projection accordingly. In [ZMPP05] Zelnik-Manor et al. proposed the use of the conformal projection Mercator and (its variant, the traverse Mercator) to project panoramas as better alternatives to the equirectangular, and central cylindrical projections (a conformal projection maintains angles at any point, preserving the shape of small objects). We implemented both projections in panotools and Hugin, and became widely used. Based upon their success we proceeded to investigate the use of other map projections in the creation of panoramas.

The main contributions of this paper are: we introduce 8 projections for the projection of panoramas. We also present a method to automatically remap photographs from its original projection to a different one.

In section 2 we propose various map projections that we believe can be used to project panoramas. In section 3 we discuss conformal projections, which we believe have large potential for aesthetic panoramas. In section 4 we propose the notion of hybrid projections to automatically combine different projections into a single panorama. In section 5 we describe a method to automatically determine the place that a photo corresponds to within the viewable sphere; with this information any photo can be automatically mapped from one projection to another. We end with conclusions and future work.

2. New projections for the Viewable-Sphere

We now proceed to describe the projections that we have found with the most potential for the creation of panoramas. Figures 2 and 3 show the same spherical panorama rendered with each of the projections and the same person appears in different areas of the panorama to facility their comparison.

2.1. Miller Cylindrical

The conformal nature of Mercator makes it very useful to display spherical and cylindrical shapes (such as people). But one of the major disadvantages of the Mercator is that it cannot present the entire sphere (it has height equal to infinity when V-FOV=180 degrees). The Miller Cylindrical projection, presented by O.M. Miller in 1942 [SV89] was designed as a variant of the Mercator to overcome this limitation. Its forward, and inverse formulae can be found

in [Sny87]. The Miller is not conformal, but it is very close to be conformal that its can be perceived (for photographic and artistic uses) as if it is was. Close to the horizon the Miller and the Mercator are almost identical. These properties make it one of the best cylindrical projections to present the entire sphere. Its aspect ratio (width/height) is approximately 1.35.

2.2. Lambert Cylindrical Equal-Area

One of the main disadvantages of cylindrical projections is that they expand the areas close to the poles. In many panoramas this space is of no particular interest (corresponding to the floor, and the nadir—either the ceiling or the sky). The main advantage of the Lambert Cylindrical Equal-Area (presented by J.H. Lambert in 1772) is that it compresses these regions. This also becomes its main disadvantage: the zenith and the nadir are usually unrecognizable. Figure 1 shows the use of this projection, in which the poles are highly compressed. We believe that it can be used to present a reasonable view of an interior room of normal height that has no special features in the ceiling or the floor. It can present the entire sphere, and its aspect ratio is 3.14, making it the most oblong of all the projections. Its forward, and inverse formulae can be found in [Sny87]. We propose the Lambert Cylindrical Equal-Area as a replacement of the Sinusoidal as an efficient storage and transmission format for panoramas due to the following reasons:



Figure 1: Lambert Cylindrical Equal Area. The poles are compressed and the emphasis is in the area in the middle. Compare this image to figure 5; the tripod used to take these photos is not recognizable. Image ©D.M. German

- Like the Sinusoidal, it is equal area (any two regions with same solid-angle use the same number of pixels).
- Its shape is rectangular, requiring no alpha channel to mask the actual image (the sinusoidal needs one).
- The rectangle that bounds the Lambert Cylindrical Equal-area is smaller than the sinusoidal's.
- Its forward and inverse formulae require one less multiplication.

2.3. Lambert Azimuthal Equal-Area

It was presented by J.H. Lambert in 1772 (at the same time as his Cylindrical Equal Area). Its main feature is that, by being equal area, it tends to use little space towards the edges (as compared to the rectilinear and the stereographic). Its aspect is a circle and it can present the entire sphere.



Miller Cylindrical



Lambert Cylindrical Equal-Area




Architectural Cylindrical



Albers Conical Equal-Area (with parameters 0 and 60)



Lambert Azimuthal Equal Area



Peirce Quincuncial

Figure 2: The same viewable sphere, mapped using each of the projections. Images ©D.M. German. Used with permission.



Guyou



Adams Hemisphere in a Square (for both hemispheres)



Adams World in a Square I and II

Figure 3: See caption of figure 2

2.4. Albers Equal Area Conical

Presented by H.C. Albers in 1805, takes two parameters (its standard parallels, corresponding to the parallels where the cone of projection is secant to the sphere). When both are one of the poles then it becomes a Lambert Equal Area Azimuthal. When both are the horizon, it becomes a Lambert Equal Area Cylindrical. We have not found any particular use for this projection, except perhaps in orientation plates. It can present the entire sphere and its aspect ratio can vary from a circle to a rectangle, depending on the choice of standard parallels.

3. Conformal Projections using complex curves

One of the most interesting periods in the creation of novel map projections was the result of advances in the complex algebra and elliptic integrals. Schwarz demonstrated that a circle can be conformally presented by the interior of a regular polygon of n sides (see [Lee76] for an extensive survey of

conformal representations of the sphere based on this principle). We proceed to implement three projections based on this principle: *Peirce Quincuncial*, *Adams World in a Square I*, and *Adams World in a Square II*. Until the advent of computers it has been extremely difficult to compute the inverse formulae for these equations (the inverse of the integral), making them impractical ([Mel31] cites this as one of their main disadvantages). For instance, when Peirce presented his projection he supplied a look-up table to assist in the calculation of the forward and the inverse (see [Pei79]). For these three projections we could not find the inverse formulae published; instead we computed them using Maple.

3.1. Peirce Quincuncial

It was the first map projection that used elliptical functions [SV89]. In its normal aspect this projection resembles an azimuthal (it is centered around the nadir). It is conformal except at 4 points where the horizon crosses the meridians at 0, 90, 180, and 360 degrees. Aside from being conformal, the Peirce Quincuncial has two important features: first, it preserves area better than the Mercator and the stereographic; second, it can be tessellated an infinite number of times. This second property has a tremendous potential for artistic purposes. The complex curves convert the straight lines in the sphere as smooth curves that are have a very pleasant periodicity. Figure 4 is an example of an sphere projected to a Peirce, tessellated 4 times and then cropped to increase its aesthetic impact. We expect the Peirce to have a similar impact to that of the stereographic in the creation of aesthetically interesting panoramas.



Figure 4: Tessellated Peirce. Image ©S. Pérez-Duarte; mapped by D.M. German. Used with permission.

3.2. Guyou and Adams Hemisphere in a Square

The Peirce can be used to create two other projections: one of its oblique forms corresponds to the Adams Hemisphere in a Square; its transverse can be used to create the Guyou projection of the hemisphere (both conformal). In their original form they were used to map a hemisphere, but they can be used for the entire sphere by tessellating the projections of each hemisphere.

3.3. Adams World in a Square I and II

Adams presented his projections in 1925 [Ada25]. They present the entire sphere in a conformal manner. The I exaggerates the size of objects along the poles but it is otherwise a good representation of the viewable sphere. The II presents the world in a square that is rotated 45 degrees. Both can be tessellated, with aesthetically interesting results. Figure 5 shows how the poles are well represented with this projection.



Figure 5: Adams World in a Square I. There is emphasis in the poles while maintaining the areas closer to the horizon. The tripod used to take these photos is clearly discernible. Compare this image to figure 1. Image ©D.M. German

3.4. False Multi point-of-view projections

Perhaps the most important feature of the Peirce, Guyou and three Adams projections is that, from the set of projections that can map the entire sphere, they provide the best balance across the entire image, including the nadir and the zenith. In a way, they can be seen as multi-point of view: they do not rely on the viewer to center the attention on any particular point (as the azimuthals do), or line (as the cylindricals do) and every area of the sphere is recognizable. We believe that the Guyou and Adams Hemisphere in a Square are good candidates to replace the equirectangular as a “thumbnail” of a panorama (usually used to link to its QTVR equivalent).

4. Hybrid projections

As we mentioned before, the choice of a projection is a compromise. There are situations in which different regions of a panorama would benefit from different projections, yet, the panorama maker has to select one of them over the other. Since the Renaissance painters have understood and exploited the advantages of combining multiple projections from the same viewpoint (for instance, Aiken describes how Massaccio masterly combined several projections in his “Trinity” fresco [Aik95]. Zorin and Barr presented a method to reduce distortion in wideangle images [ZB95]. Zelnik-Manor et al. proposed the use of different viewpoints to minimize distortion in wideangle panoramas [ZMPP05] (we have found that this approach creates abrupt changes in straight lines, which yields “false” views: it can create believable corners where non-existed, or remove them where they were, potentially fooling the viewer; furthermore it requires user interaction to choose the best seam lines).

Furthermore, using several projections from the same viewpoint (or center of projection) in the same image should not be confused with using different viewpoints, as the two concepts are orthogonal to each other: an skilled painter can use both in the same painting. It could be argued that a skilled operator of an image manipulation system could combine different projections using masking & compositing, the way an artist is able to combine them in a painting (this will be very time consuming). We investigated if it was feasible to use two or more projections in the same panorama in an automatic way. We observed that the main challenge of using two projections is to provide an “invisible” seam where the projections merge. This could be done in two ways: a) by creating a transition area (as [ZB95] does it for wideangles), where one projections slowly transforms into the other, or b) by finding projections that perfectly align at one plane of the sphere without abrupt changes in the direction of lines that cross from one projection to the other (a full spherical panorama requires a plane, but a partial panorama requires only a segment of a plane that crosses its are of interest). This section explores the latter approach.

4.1. Hybrid Cylindrical

We observed that the horizon creates a plane along which cylindrical projections can be easily combined. The formula to map the x-coordinate of all the cylindricals is the same, and the formula to map the y-coordinate is very similar for values close to the horizon. When any two cylindrical projection are combined, one for the top hemisphere, and a different one for the bottom hemisphere, a very smooth area of transition along the horizon is created.

4.2. Architectural Cylindrical

When we are in an open space, particularly one with buildings around, or in an atrium, our sight tends to move up-

wards. Our attention is focused in the top hemisphere rather than the lower area of the bottom hemisphere. Unfortunately a projection that puts more emphasis on the region around the zenith (such as the Mercator or the Miller) also places a similar emphasis in the opposite direction. In the real world, as we perceive the sphere around us, the nadir is occupied by ourselves, and we usually don’t explore the area immediately around us (except in situations where the floor is of particular interest). One combination that we have found useful for this situation is the is Miller (for the top hemisphere) and Lambert Equal Area (for the bottom one). We have named this projection *Architectural Cylindrical*. Its aspect ratio is 1.92. An interesting side effect of this projection is that the horizon is shifted from exactly the middle, to 30% from the bottom, which is close enough to follow the rule of thirds, as shown in figure 6.

There are, however, situations in which the opposite is true: the top hemisphere is less interesting than the bottom, such as relatively flat landscapes where the top hemisphere is occupied by the sky. In this case our sight tends to explore the bottom hemisphere. For this reason the Architectural projection takes a parameter indicating the orientation of the projection: +1 (default, uses Miller for top hemisphere), and -1 (for the opposite). Its main disadvantage is that some objects might appear distorted if the extend a significant V-FOV and cross the horizon plane.

Figure 7 shows the interior of the Reims Cathedral using the architectural projection. Figure 8 shows the same panorama using the equirectangular projection.



Figure 6: Architectural projection. The horizon has been shifted from the middle to almost 30% from the bottom of the image. Image ©D.M. German. Used with permission.

5. Automatic remapping of photographs

In order to be used as intended, some projections require the horizon to be placed properly (such as the cylindricals), while other projections are not very affected by its location (such as the azimuthals). This property of the azimuthals permits the automatic remapping on one image from one azimuthal projection to another. This method is commonly used to automatically map the image casted by a fisheye lens (either equidistant, or equal-area) into a rectilinear or



Figure 7: Architectural projection, which highlight the detail in the ceiling more than the floor. Compare this image with figure 8. Image ©S. Pérez-Duarte



Figure 8: Equirectangular projection of the image in figure 7. Equirectangular is the most commonly used projection to display full spherical panoramas. The floor occupies a large area, but the ceiling is compressed. Image ©S. Pérez-Duarte

stereographic image (stereographic fisheye lenses—which are conformal—are perceived to create a more pleasant image than fisheyes, but are difficult to manufacture) and is known as “defishing”. Several software applications are available in the market to perform this transformation automatically. The fisheye image is automatically “defished” into a rectilinear projection where the center of projection is the same in both.

In order to defish a photo with a projection where the position of the horizon is important, one would need to have a perfectly leveled photograph: the camera is perfectly aligned with the horizon (its pitch and roll are both zero). The pitch is the angle of elevation from the horizontal axis of the camera; the roll is the angle of rotation around the optical axis. When the camera is not perfectly leveled there are two approaches: one can specify the pitch and roll manually, or one can identify vertical or horizontal control lines as a reference. In the latter case image registration software is used to compute the pitch and roll from these input control lines. Knowing the pitch and roll is crucial because they indicate what part of the sphere the photograph corresponds to, in order to properly represented the horizon in the output projection (without specifying the horizon location an image will be mapped into an oblique version of a projection, with verti-

cal lines appearing curved). Once the image has been placed with respect to the sphere, it can be properly mapped into any projection of normal aspect. Either approach (making sure the pitch and roll are zero when the photo is taken; or manually setting control lines) is time consuming.

The success of the Wii remote (WiiR) from Nintendo as an input device to measure both acceleration and inclination prompted a question: would it be possible to use an inclinometer to record pitch and roll in order to automatically map a photograph from one projection to another with an acceptable result?

5.1. Using accelerometers to record pitch and roll of a photograph

The WiiR contains three accelerometers orthogonal to each other that measure acceleration in each of the three dimensions. By combining their readings one can determine the acceleration vector in which the WiiR is moving. The WiiR can be used also to measure inclination. If the WiiR is in a stationary state (it is not accelerating) then the difference in the reading from each of its accelerometers can be used to compute its pitch and roll.

As an inclinometer the WiiR has very low precision. Each accelerometer can only reliable measure angles between -45 to 45 degrees. By combining the readings of the three accelerometers one can read pitch and roll from -90 to +90 in two planes. We wrote a program to query the WiiR sensors, then calculate, timestamp and log the WiiR pitch and roll (we modified *DarwiinRemote* for this purpose <http://blog.hiroaki.jp/2006/12/000433.html>). We then attached the WiiR to a camera via its tripod socket. The WiiR is a Bluetooth device, hence it did not require a physical link between the camera and the computer, allowing hand-held photo shooting. We synchronized the clocks of both the camera and the computer. The camera timestamps a photo with a resolution of one second. To guarantee that both the camera and the WiiR logs were synchronized every session a photo of the computer display (with the logging tool in the foreground) was taken at the beginning of a session; this photo showed the time according to the computer and it was timestamped by the camera. This information was found to be enough to synchronize the logs.

The next challenge was to determine if the photographer could keep the WiiR and the camera in a stationary state long enough to allow recording of pitch and roll. This is not a problem with the use of a tripod (the camera is always in a stationary state). Fortunately, when a photo is taken handheld the photographer is expected to keep the camera steady (there are multiple exceptions to this rule, such as photographing sport events or using panning—where the camera is rotated at a constant speed while a photograph is taken). We found that in general the camera remains in a state such that it was possible to measure its pitch and roll.



Figure 9: The Wii Remote mounted under a Canon 20d camera using the tripod socket.

Our experiments indicate that the values returned by the accelerometers vary from -20 to $+20$ (once we have normalized the value 0 to the horizon), corresponding to -45 to 45 degrees; the values are evenly spaced), giving a maximum theoretical precision of approximately 2.25 ± 1.12 degrees.

Unfortunately the WiiR is not able to measure the orientation according to the Earth pole in which the photo is taken (the yaw—the WiiR relies on an infrared emitter to determine yaw in front of a TV set).

Without yaw, it is impossible to properly place an image within the viewable sphere. If we restrict $yaw = 0$ then we can place the center of an image along an imaginary vertical line in front of the viewer. We can therefore map the image with a given $\langle yaw = 0, pitch, roll \rangle$ to its location in the viewable sphere, and from there, map it back to a flat representation (either in its original projection or a different one).

There are three practical uses for this approach: 1) **Correcting tilted horizons:** once the roll is known, it is trivial to “straighten” a photo; 2) **Correcting perspective errors:** using panoramic software (such as panotools) it is possible to automatically remap the photo taking into account the angle in which the camera was tilted-up (the pitch). For an example of this type of correction see figure 11 (specialized lenses are able to do this optically but unfortunately they are expensive, require skill to be used effectively, and can only fix parallax when the angle is small); and 3) **Remapping to another projection:** once the pitch and roll is known a photo can be mapped to any other normal projection. See 12 for an example of a remapping of a fisheye-lens image onto a Lambert Equal-Area projection.

Our experiments show that the precision of cheap accelerometers (as found in the WiiR) is enough to effectively perform all three approaches if the photo is not taken with a large telephoto (the longer the lens, the more the error is magnified). These preliminary results are promising, but it is necessary to conduct a formal experiment to properly verify them.



Figure 11: Correcting perspective. The corrected image preserves parallel lines as parallel. (the original is above)



Figure 12: Automatically remapping an image from a Fish-eye lens to Lambert Equal Area (the original is above).

With a fast enough processor these corrections could happen in the camera, in real time, and the user could preview them using the camera’s LCD display. These might be features that one day become common in our digital cameras. The ability to record pitch and roll by the camera has also important legal implications. Our brains rely on visual clues to determine where the horizon is, and what corresponds to vertical [LT04]. When those clues are manipulated, an image might be interpreted by its viewer differently from the reality that it corresponds to. We also envision a future in which a camera will record GPS coordinates (this is already possible), pitch and roll and orientation with respect to the North




Figure 10: Block diagram showing the automatic correction of images with the WiiR.

pole. This way every photo will record the place in the world where it was taken, and the orientation of the camera (in 3 dimensions). This will have a dramatic effect in the way that we use photographs for the reconstruction of spaces.

6. Conclusions and Future Work

In this paper we have proposed the use of several map projections for the presentation of flat panoramas. We have also presented the idea of hybrid projections, in which different regions of the viewable sphere are presented with different projections. Finally we propose an method based on inclinometers attached to the camera to automatically project a photograph to the viewable sphere, and then into any panoramic projection.

We believe that the use of new projections can have important uses not artistic panorama making, and in other more practical uses. Immersive systems, for example, should explore the use of cylindrical projections instead of rectilinear or fisheye for large H-FOV because they preserve vertical lines as vertical ([Hag91] listed it as an important feature of this type of systems). Also, empirical studies that evaluate the effectiveness of different projections (for a particular goal) needs to be performed. We believe that conformal mappings have a great potential for practical and aesthetic panoramas, and they need to be further explored. Hybrid projections is also a promising area. For example, an immersive system can switch from one projection to another as the FOV changes, but this transition needs to be seamless.

7. Acknowledgments

We want to thank Sébastien Pérez-Duarte, the users of the *PanotoolsNG* mailing list, and the members of the Flickr group *Lords of the Photo-stitching*, and reviewers of this document for their valuable input in the preparation of this document.

References

[Ada25] ADAMS O. S.: *Elliptic Functions Applied to Conformal World Maps*. Tech. Rep. Special publication

No 112, Department of Commerce, U.S. Cost and Geodetic Survey, 1925.

[Aik95] AIKEN J. A.: The Perspective Construction of Masaccio's "Trinity" Fresco and Medieval Astronomical Graphics. *Artibus et Historiae* 16, 31 (1995), 171–187.

[Hag91] HAGEN M. A.: How to make a visually realistic 3D display. *SIGGRAPH Comput. Graph.* 25, 2 (1991), 76–83.

[Kub86] KUBOVY M.: *The Psychology of Perspective and Renaissance Art*. Cambridge University Press, 1986.

[Lee76] LEE: Conformal Projections Based On Jacobian Elliptic Functions. *Cartographica* 13, 1 (1976), 67–101. Monograph 16.

[LT04] LIU B., TODD J. T.: Perceptual biases in the interpretation of 3D shape from shading. *Vision Research* 44, 18 (Aug. 2004), 2135–2145.

[Mel31] MELLUISH R.: *The Mathematics of Map Projections*. Cambridge University Press, 1931.

[Pei79] PEIRCE C. S.: A Quincuncial Projection of the Sphere. *American Journal of Mathematics* 2, 4 (Dec 1879), 394–396.

[Sny87] SNYDER J. P.: *Map Projections—a working manual, Professional Paper 1395*. U.S. Geological Survey, 1987.

[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.

[ZB95] ZORIN D., BARR A. H.: Correction of geometric perceptual distortions in pictures. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1995), ACM Press, pp. 257–264.

[ZMPP05] ZELNIK-MANOR L., PETERS G., PERONA P.: Squaring the Circle in Panoramas. In *Tenth IEEE International Conference on Computer Vision* (2005).