

Theory of the “No-Parallax” Point in Panorama Photography

Version 1.0, February 6, 2006

Rik Littlefield (rj.littlefield@computer.org)

Abstract

In using an ordinary camera to make panorama photographs, there is a special “no-parallax point” around which the camera must be rotated in order to keep foreground and background points lined up perfectly in overlapping frames. Arguments about the location of this point are both frequent and confusing. The purpose of this article is to eliminate the confusion. It explains in simple terms that the aperture determines the perspective of an image by selecting the light rays that form it. Therefore the center of perspective and the no-parallax point are located at the apparent position of the aperture, called the “entrance pupil”. Contrary to intuition, this point can be moved by modifying just the aperture, while leaving all refracting lens elements and the sensor in the same place. Likewise, the angle of view must be measured around the no-parallax point and thus the angle of view is affected by the aperture location, not just by the lens and sensor as is commonly believed. Physical vignetting may cause the entrance pupil to move, depending on angle, and should be avoided for perfect stitches. In general, the no-parallax point is different from the “nodal point” of optical designers, despite frequent use (misuse) of that term among panorama photographers.

Introduction

In using an ordinary camera to make panorama photographs, there is a special “no-parallax point”¹ around which the camera must be rotated in order to keep foreground and background points lined up perfectly in overlapping frames. Arguments about the location of this point are both frequent and confusing. They involve statements like these:

“The light rays that form an image must all go through the entrance pupil.”

“Adding an aperture does not change the path of light through a lens.”

“The angle of view is determined simply by sensor size and lens focal length.”

These apparently simple statements, and others like them, have been the source of great confusion about the location of the no-parallax point. (The third of these statements is even wrong, in the general case!)

The purpose of this article is to eliminate the confusion. It explains in simple terms why the no-parallax point is located at the entrance pupil, is also the center of perspective, and can be moved by modifying just the aperture, while leaving all refracting lens elements and the sensor in the same place. It also explains how the angle of view (Panorama Tools *hfov*) is actually determined by the position of the no-parallax point, not just the lens and sensor. And finally, it explains why physical vignetting can cause bad stitches in addition

¹ See the PanoTools Wiki, <http://wiki.panotools.org>.

to dark bands. Along the way, it clarifies that the no-parallax point is far different from the lens designer's "nodal point," despite widespread use of that term among panorama photographers.

Caution: to fully understand the no-parallax point, you may have to give up some long-held beliefs. For example, if you insist that "angle of view is determined simply by sensor size and lens focal length", then you will be in trouble. That statement is true only in specific cases that unfortunately dominate the literature and have become entrenched in our thinking. Just keep an open mind -- all will be explained.

Organization of the Article

First, we summarize in words why there has been debate over the position of the no-parallax point. Second, we use basic ray-tracing and a thin-lens model to analyze exactly what happens under several conditions that historically have been confusing. Several photographs are provided to illustrate and confirm this analysis. Third, we illustrate how the analysis applies to actual multi-element lenses, which allows us to illustrate and explain that the no-parallax point is generally not at what optical designers would call a "nodal point". Fourth, we study the concept of "angle of view" as it relates to appearances of the images and to settings of the Panorama Tools lens parameters. Finally, we explain how vignetting relates to all this.

The Debate

There is a simple argument why the no-parallax point must be at the aperture, or more precisely, at the **"entrance pupil"**, which is where the aperture appears to be as viewed through the front of the lens. The argument goes like this. Imagine two object points positioned so that they are perfectly lined up, one behind the other, as seen by the camera. A ray of light passing from the background object point straight through the foreground object point must go through the aperture; otherwise it could not form the image. If the aperture is moved so that that ray cannot get to the sensor, then the two points will not appear to be in line. Therefore the entrance pupil must stay in line with that ray. It must also stay in line with corresponding rays cast through all other pairs of points that are lined up. The only way to accomplish that is to rotate the camera and lens around the entrance pupil. Similarly, the center of perspective lies at the intersection of straight lines between all the points that appear to line up, so the center of perspective must coincide with the no-parallax point. Therefore the no-parallax point, the center of perspective, and the entrance pupil must all be at the same place. Case closed ... or so it would seem.

Unfortunately, there is another simple argument why the no-parallax point could not be determined by the aperture and should be determined by something else. That argument goes like this. The location of the no-parallax point – what lines up with what – is determined by the path of the light rays. But the path of light rays through a lens is determined only by the refractive elements, not by the aperture. Therefore the location of the no-parallax point cannot be determined by the aperture, but must be an intrinsic property of the lens.

In addition, there is a third argument, more complex and technical, about why the no-parallax point should be at the “front nodal point” of a lens. This argument has to do with the preservation of angles for light rays that take particularly simple paths through the lens. Space prevents even summarizing the argument here, but perhaps you can take our word that it sounds convincing by itself.

Obviously, not all of these arguments can be correct. Hence the debate: which one is?

Many experiments were performed, but the results were often more confusing than helpful. Careful measurement of a few well-behaved lenses indicated that their no-parallax points were at least very close to the location of the entrance pupil, and distinct from the optical nodal points. Additional experiments, adding an external aperture, indicated that the no-parallax point shifted to very near the external aperture, but perhaps not quite at it. Still more experiments determined that it was possible to set up conditions with an external aperture such that the no-parallax point appeared to move smoothly over a distance of several inches as the built-in lens aperture was changed from fully open to fully closed.

Taken together, all of this material can be nicely summarized in one word: Aarrggh!!!

Ray-Tracing a Thin Lens

It turns out that great insight can be obtained by carefully ray-tracing an extremely simple optical system: a single “thin lens”.

To trace rays through a thin lens, one needs only three rules:

1. All rays entering the lens parallel to the optical axis are bent to pass through a single point on the optical axis, exactly one focal length away from the plane of the lens.
2. All rays passing through the center of the lens are not bent at all.
3. All rays starting from a single object point are bent to pass through a single image point, regardless of where they enter the lens.

These rules are completely sufficient to generate the standard algebraic model of a thin lens, that $1/f = 1/o + 1/i$, where f is the focal length, o is the distance to the object, and i the distance to its image. However, for the purposes of this paper, it is more effective to use the rays directly.

Consider the situation shown in Figures 1 and 2. On the left side, two arrows serve as objects. On the right side, there is an image plane, designated by IP. The right arrow happens to be in focus at IP; the left arrow focuses somewhat closer to the lens and is out of focus at IP. (For clarity, a very wide lens is shown – about $f/0.5$! This does not affect the analysis.)




Figure 1. Ray Tracing a Thin Lens

Figure 1a shows only the simplest rays – the ones determined by rules 1 and 2 above. Once a couple of these rays are placed, it is straightforward to generate any number of other rays that might be needed to analyze various situations.

Now consider figure 1b, which shows the other rays – all of them – that we will need for this discussion.

Notice that there is a fan of rays for each object point, for example the red rays passing through the head of the left arrow. Figure 1b shows the maximum range of available rays. Adding an aperture anywhere in the system will clip away some of these rays, so that the width and location of each fan depends on the aperture.

For object points that focus perfectly at IP, the width and location of the ray fan is of no importance. All of these rays converge on the same image point. Clipping away some of them has no effect except to make the image less bright.

For object points that do not focus perfectly at IP, the width and location of the ray fan is of critical importance. These rays form a blur in the image plane at IP. Viewers identify the center of the blur as the apparent location of the object point. When the lens is wide open, the blur is large. When some of the rays are clipped (blocked) by an aperture, the remaining rays form a smaller blur, and in addition the center of the blur may be moved. If the clipped fan becomes sufficiently small, the image of an out-of-focus point may appear sharp, but still its location may change depending on which rays get through the aperture.

Figure 2 shows what happens when a small aperture is introduced at various positions. The aperture, shown by a yellow circle, clips away all rays except the ones that pass through it. For clarity, these diagrams show only the rays for the out-of-focus arrow, and the slightly blurred image of that arrow at IP is shown as a red bar.




Figure 2. Ray Tracing a Thin Lens Plus Small Aperture

In figures 2a, 2b, and 2c, we leave the aperture in the plane of the lens, but shift it cross-axis. In 2d, we move the aperture in front of the lens, nearer to the objects.

As the aperture shifts, the in-focus image of the right arrow stays in exactly the same place, because all the rays for the right arrow converge on a single point in the image. But the position and size of the out-of-focus image changes depending on the aperture position.

In 2a, 2b, and 2c, the out-of-focus image just shifts up and down. Notice that in 2a the images of the arrowheads are aligned, in 2b the centers of the arrows are aligned, and in 2c the tails of the arrows are aligned. In each case, the image at IP is consistent with the center of perspective being at the location of the aperture.

Figure 2d is perhaps most interesting. The out-of-focus image has become smaller! Again, this is consistent with the center of perspective being located at the aperture, now relatively closer to the in-focus arrow than the out-of-focus.

Summarizing so far... In general, when we impose a small aperture, the in-focus image stays the same, except for getting dimmer as we make the aperture smaller. The out-of-focus image also gets dimmer, by the same amount on average, but this is accomplished by leaving intact the portion of the blur that corresponds to having the center of perspective at the aperture, while eliminating all other portions of the blur. Not only do the blur circles get smaller, but their centers shift location. If you have objects at different distances, then their images shift and scale by different amounts that are appropriate to having the center of perspective at the aperture. All this happens by just clipping the ray fans.

An equivalent summary is just this:

The unstopped lens would collect rays representing many different centers of perspective. The aperture selects a subset of those, and the effective center of perspective belongs to the rays that were selected.

Photographic Demonstration

Diagrams are very nice, but they are easy to misinterpret and to draw incorrectly. Perhaps a direct photographic demonstration will be helpful.

Figure 3 show photographs taken with a close-up setup to make it easy to see effects like those described above.

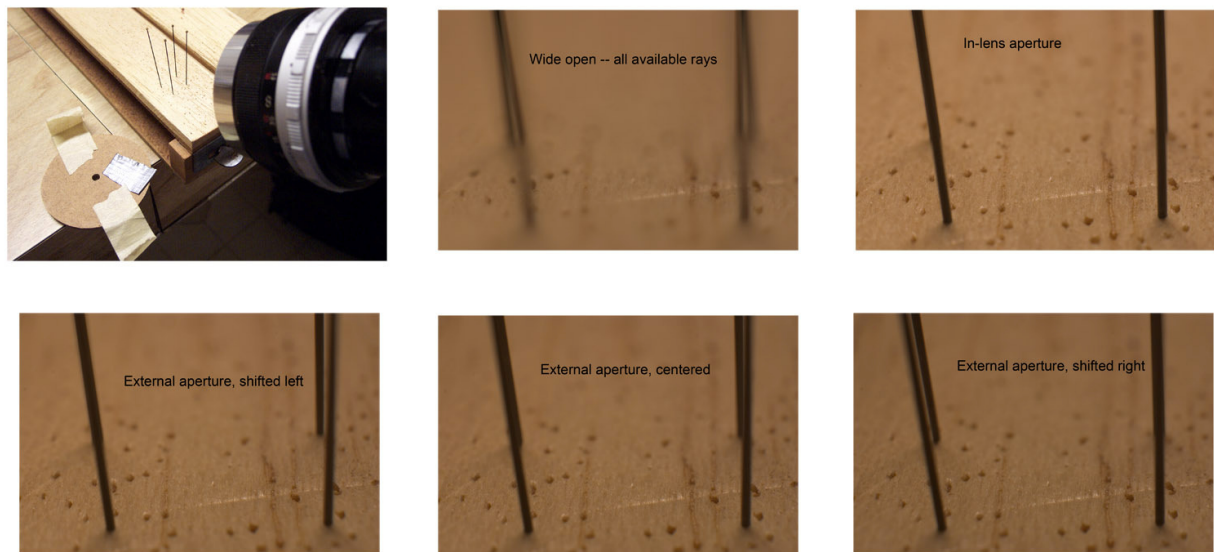


Figure 3. Experimental demonstration of shifting perspective by moving the aperture.

The setup involves 4 pins stuck into a piece of balsa, photographed with a Canon 300D through a 55mm prime lens on extension tubes. Pictures were taken using 5 different aperture conditions, with NO changes to focus, lens, camera, or subject positions.

Figure 3, top center, shows the picture with aperture wide open. This represents all available rays.

Figure 3, top right, shows the picture with the in-lens aperture stopped clear down. This puts the entrance pupil in the middle of the lens. The pins are carefully placed so that they line up in this case.

Figure 3, bottom row, shows the picture with the in-lens aperture wide open, and an external aperture taped to the front of the lens. The three pictures show what happens

when the aperture is shifted slightly left, centered, or slightly right.

Notice in Figure 3, that each of the pictures with a small aperture has a different perspective. In each case, the in-focus plane is not affected by the aperture, but the images of out-of-focus points are shifted and scaled by different amounts depending on their distance from the entrance pupil, so that the center of perspective remains at the entrance pupil.

These same effects are shown more vividly at <http://www.janrik.net/PanoPostings/NoParallaxPoint/ApertureStack.gif>, which is an animated gif that cycles between images taken with the internal and external apertures.

Real Lenses (Thick Lens Model)

The previous analysis for a thin lens should make clear that the no-parallax point and the center of perspective are both located at the aperture, at least when the aperture is located in the plane of the lens or on the object side of it.

But what happens with a real lens, where the aperture is located between the lens elements? And just what is this thing called a “nodal point”, anyway?

We do not have space enough here to discuss these issues in full detail, but a brief summary should answer the questions well enough for our purposes.

When the aperture is located anywhere behind the front lens element, and you look through the front of the lens, the aperture simply appears to be at a different location and size than it really is. (Think of looking at a ring through a magnifying glass.) This apparent position and size of the aperture is called the “**entrance pupil**”.

Thinking again about the earlier discussion, it should now be obvious that the no-parallax point and the center of perspective are actually located at the entrance pupil. Rays of light directed toward the entrance pupil are bent by lens elements, pass through the aperture, are bent more by other lens elements, and eventually hit the sensor. The details of the bending are quite irrelevant for determining the center of perspective.

It turns out that the details of the bending are quite irrelevant for many other purposes also. In fact, most lenses can be approximated quite closely by a simple “thick lens model” that has two refracting surfaces separated by a gap. Rays can be traced through a thick lens using rules that are similar to those of a thin lens, but with more parameters. Where a thin lens has only one parameter (focal length), a thick lens has six. Under most circumstances, some of the parameters become redundant.² The thick lens model then has three independent parameters: one focal length and two “focal points”, from which two

²When the refractive index of the medium is the same on both sides of the thick lens, then the front and rear focal lengths have the same magnitude and the nodal points coincide with the principal points. This is true all the time with traditional cameras. It is not true, for example, with oil immersion microscope objectives.

“nodal points” = “principal points” are determined. Planes perpendicular to the optical axis and passing through the principal points are called “principal planes”. These few parameters tell you everything you need to know for tracing rays through a thick lens. (More precisely, the rules tell you how to construct rays around a thick lens – they do not accurately show the paths of rays inside the lens).

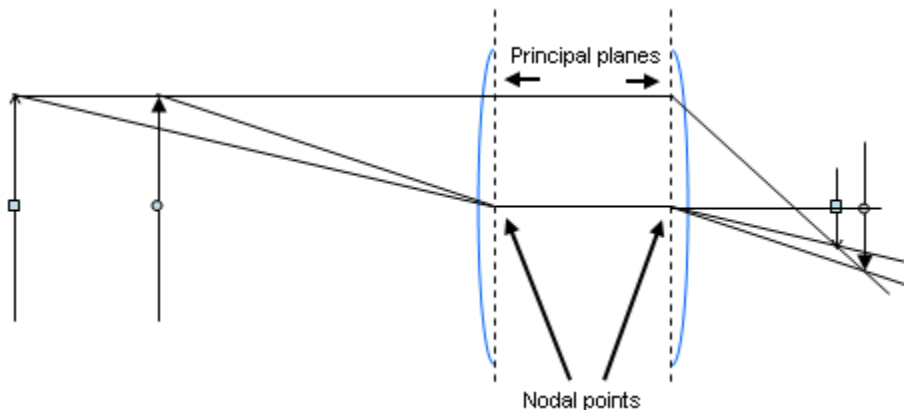
For general rays through a thick lens, the rules are fairly complicated. However, for two kinds of rays, the rules are very simple:

1. All rays passing through one principal plane parallel to the optical axis are pretended to bend at the other principal plane so that they pass through that plane’s focal point.
2. All rays directed toward one nodal point emerge as if coming from the other nodal point at the same angle.

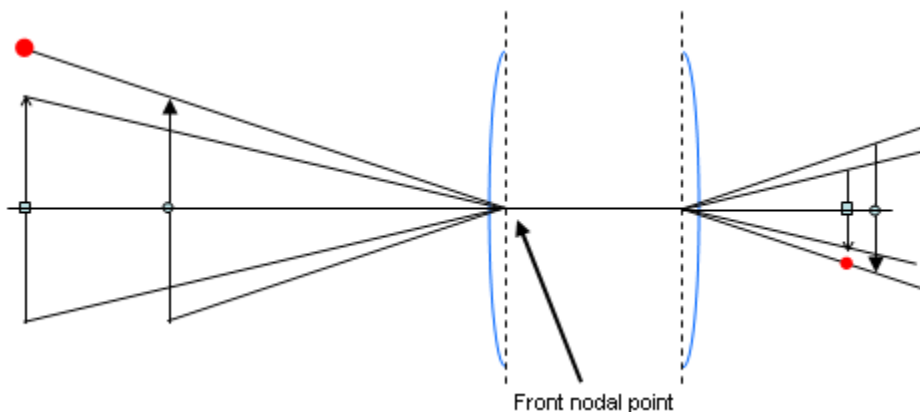
In addition, we have the same third rule as for a thin lens:

3. All rays starting from a single object point are bent to pass through a single image point, regardless of where they enter the lens.

Because of the simplicity of these rules, they are often used to construct diagrams. For example, in our case with two arrows and a thick lens (instead of thin), a few of the rays look like this:

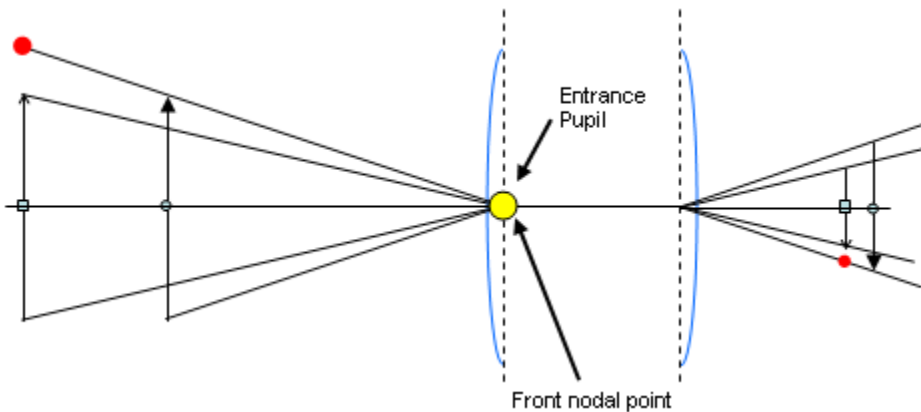


From here, it is natural to draw another even more suggestive diagram:




We believe that diagrams like this last one have contributed much to the confusion about the no-parallax point. There is no debate that these rays are correctly drawn, and faced with this diagram, it seems obvious that we must rotate around the front nodal point. What other point could we rotate around, that would keep the red dot and the black arrowhead lined up with each other?

The problem, of course, is that the way this diagram is drawn forces an incorrect conclusion. By choosing to draw these rays, and only these rays, we have drawn the diagram as if there were a very small entrance pupil centered on the nodal point. The result should be no surprise by now – placing the entrance pupil at the nodal point puts the center of perspective there too.



In real lenses, there is no particular relationship between locations of the entrance pupil and the front nodal point. It is quite likely that rays directed at the front nodal point will not even get through the lens. Images are formed by rays directed toward the entrance pupil, which are bent to pass through the physical aperture.

This diagram³ should make these points clear.



The distinction between front nodal point and entrance pupil = no-parallax point = center of perspective is not just academic. These locations can be very far apart in practical lenses, such as the telephoto shown here.



Primary Nodal Point and Entrance Pupil for the Pentax Super-Takumar 200mm f/4 lens

In passing, it is important to note that the rear nodal point does play a legitimate role in panorama photography, but only with special swinging lens cameras. In those cameras, the film is wrapped into a circular arc behind the lens. The camera stays in one place while the lens is rotated in front of it, the film being exposed piecemeal by a slit shutter that tracks the lens. In such cameras, it is important that the image as a whole be

³ Basic diagram and ray tracing by LINOS Photonics WinLens (<http://www.winlens.de/>). Annotation added.

stabilized on the film, and this is accomplished for objects at infinity by rotating around the rear nodal point. If the entrance pupil is somewhere else (as it is in most cases), then the center of perspective does move as the lens rotates. However, this does not introduce visible parallax because the lens rotates only a very slight amount while the shutter exposes each spot of film.

Angle of View and Other Lens Parameters

Now we want to address how to describe images in terms of their “angle of view” and other parameters that are needed to correctly interpret them and turn them into a panorama.

Correctly interpreting the geometry of images is a general issue in photogrammetry and has been addressed in some detail for stereoscopy⁴. However, in this article we will discuss it from the standpoint of panorama generation, and in particular the parameters needed by the Panorama Tools software.

Caution: Please stop here and be sure that you are comfortable with the discussion in previous sections of this article. If not, then please go back and study it until you understand it well. You need to be clear that:

- The refractive elements of the lens form ray fans that are clipped by the aperture.
- The image of in-focus points is not affected by the aperture.
- The images of out-of-focus points are scaled and shifted by clipping the ray fans, so that the resulting image has its center of perspective at the center of the entrance pupil.

When these points are clear, continue reading.

Now suppose, for the sake of discussion, that you have set up the close-up gear used for Figure 3, complete with the added aperture in front of the lens, and you want to use that gear to shoot a mini-panorama. How would you do it?

(We admit, actually doing such a thing is arguably crazy. But it’s an excellent model problem. If you can get the right result in this case, there’s a good chance that you really do understand the no-parallax point and can get the correct result in other weird cases too. Trust us, working through this one definitely helped to flesh out our understanding.)

So, how would you shoot the mini-panorama?

It’s probably clear by this point that you need to rotate the camera around that external aperture, which is acting as the entrance pupil and determines the center of perspective. But rotating around the proper point is only part of the problem. You must also provide values for what Panorama Tools calls “lens parameters”, notably the “horizontal field of view” (*hfov*) and a couple of “lens offsets” (*d* and *e*). What about those?

⁴ See Bercovitz, <http://www.angelfire.com/ca2/tech3d/images/persp.pdf>.

First, let's be clear that when Panorama Tools speaks of "horizontal field of view" (*hfov*), it really means "horizontal angle of view". Angles have a vertex, of course, and this prompts a surprisingly interesting question:

Where is the vertex for the angle of view?

This question is more subtle and derailing than it probably appears at first. It gets at the very basic issue of what "angle of view" means.

The literature is full of discussions about "angle of view", but hardly any of them precisely define what they mean.⁵ Instead, they draw simple diagrams and quote formulas derived from those diagrams. This presents a problem for understanding the more subtle issues discussed in this article. The problem is that the standard simple diagrams and formulas are only approximations. They are fairly accurate most of the time, obviously inaccurate some times, and wildly wrong once in a while, but they seem so obvious and have been repeated so often that many people accept them without question. If you are one of those people, then please be aware of that fact and prepare to do some mental stretching.

To begin the discussion, let's consider a problem that the standard diagrams do not address.

Suppose that someone gives you an image of a three-dimensional scene and asks you to determine its angle of view. No information is provided about the camera or lens, but you do have access to the scene where the image was taken. Is the angle of view well defined? How would you determine it?

The answer to the first question is "Yes, the angle of view is well defined." For the second question, the method of determining the angle of view is probably obvious: find where the camera was positioned, and from that point, measure the angle between object points that appear at the edges of the image.

How do you find where the camera was positioned? Easy – you just move around until you see foreground/background points lined up in the scene the same way they are in the image. Wherever your eye is when that happens, that's where the camera was positioned, and that's the point around which you measure the angle-of-view.

But remember, the image was formed by light rays that were directed toward the entrance pupil. So what you have actually found, in carrying out this exercise of lining up

⁵ Actually, in preparing this article we could not find any that do precisely define what they mean. The concept "angle of view" seems to be treated as obvious. It isn't.


foreground and background points, is the location of the entrance pupil when the image was formed. This result is critical, so we will emphasize it:

The angle of view is the angle between object points that appear on opposite edges of the image, as measured around the entrance pupil (the center of perspective, the no-parallax point).

This statement by itself may not bother you very much, but it has consequences that you may find disturbing. The angle of view is determined by the position of the entrance pupil, and the position of the entrance pupil is determined by aperture placement as well as other factors. This means that **moving the aperture changes the angle of view of the system**, even though there has been no movement of the subject, sensor, or refractive elements of the lens.

Consider again Figure 3. If we did not tell you how the pictures were taken, you would say that the lower pictures were taken with a shorter lens, positioned closer to the pins, and having a wider angle of view. In fact the lower pictures cannot be distinguished from pictures that actually were shot with a shorter lens having a wider angle of view -- because exactly the same rays are captured.

For clarity, let us now discuss why this result differs from most presentations that appear in the literature. Such presentations are easily found – just ask any search engine to find Web pages containing *lens “angle of view”*.⁶ If you can find a diagram at all, it generally looks like this:



To go with the diagram (or more frequently by itself), a standard formula will appear:

$$\text{angle-of-view} = 2 \arctan\left(\frac{d}{2f}\right) \text{ for sensor dimension } d \text{ and focal length } f$$


Somewhat less often, the discussions add that this formula is valid only when the lens is focused on a subject at infinity, and that for close-up and macro work, the formula must be adjusted to account for the closer focus, say

$$\text{angle-of-view} = 2 \arctan\left(\frac{d}{2i}\right) \text{ with image distance } i = f(1+m) \text{ for magnification } m$$

⁶ See, for example, http://bobatkings.com/photography/technical/field_of_view.html or http://en.wikipedia.org/wiki/Angle_of_view .

By this point, we hope that you (the reader) can see the conceptual shortcoming in this diagram and these formulas: **they assume a thin lens model with the aperture located in the plane of the lens.** If the real lens and the real aperture behave differently from this model, then the diagram will not be appropriate and the formulas will be wrong.

We have already shown a couple of diagrams that come close to illustrating this point. Let us repeat them with slight modifications to make the situation clear.




These two diagrams correspond to figures 2b and 2d shown earlier, but here we emphasize the angle of view subtended by the solid arrow, which fills the sensor. Notice that both diagrams show the same object, sensor, and lens, with everything in exactly the same place except for the aperture. And yet, on the left diagram the angle of view α_1 is about 37 degrees, while in the right diagram the angle of view α_2 has increased to about 68 degrees.

If you are having trouble making sense of this, remember the key points made earlier. Moving the aperture does not change the image size for object points in the plane of focus. By clipping their ray fans, the aperture does change the image size for object points in front of or behind the plane of focus. The resulting image always has its center of perspective at the entrance pupil. If the entrance pupil moves closer to an object in the focus plane, then the angle subtended by that object increases. When that object happens to fill the image, it becomes obvious that the angle of view of the image has also increased. Moving the aperture does not change the field of view; it does change the angle of view.

One final diagram will illustrate how badly wrong the standard formula for angle-of-view can be. If we place the aperture at the rear focal point for a thin lens, then we force the system to be “telecentric” on the object side, with each cone of rays being oriented

perpendicular to the object plane.



The properties of such an optical system are interesting and useful. For example, lenses that are telecentric on the object side are highly valued for machine vision and optical gauging because they exhibit constant magnification – the image of an object that is slightly out of focus is merely blurred, not changed in size. However, this implies that the entrance pupil is at infinity, and the corresponding angle of view is zero! To understand this, it may help to think of imaging the object with a series of longer and longer telephoto lenses, increasing the entrance-pupil-to-subject distance as needed to keep the same image size. As you do that, you get closer and closer to infinity perspective, while the angle of view gets smaller and smaller. As far as the perspective and angle of view are concerned, a telecentric lens simply uses clever optics to carry this process to its limit, pushing the entrance pupil to infinity while the lens itself remains fairly small.

We hope that this lengthy discussion has made the point clear: the standard formulas for angle of view are not precisely correct and can be significantly wrong. For precise stitching, the angle of view (Panorama Tools *hfov*) must be measured around the entrance pupil, the center of perspective, the no-parallax point. Measuring around any other point will produce stitching errors.

Now about the shifts, d and e . These parameters have two common uses. First, they establish the center of distortion for fisheye lenses and for the $a/b/c$ polynomial corrections of all lenses. Second, they can be used for “flat stitching”, by locking $\text{pitch}=\text{yaw}=0$ and using d and e to correct for relative movement between the subject and the sensor. Off-center apertures present a third use for d and e .

Within the geometric model implemented by Panorama Tools, the shifts d and e actually encode part of the information about where the center of perspective is located. In more detail, the center of perspective is assumed to reside on a line perpendicular to the sensor, passing through image coordinates $(\text{width}/2+d, \text{height}/2+e)$, and located at a distance away from the sensor that is sufficient to make a centered image subtend the angle $hfov$. As noted by Bercovitz:

“The distance from an image to its correct perspective point is numerically equal to the magnification of the in-focus object at the image plane times the distance from the entrance pupil of the lens to the in-focus object.”⁷

Let b equal this distance from an image to its correct perspective point. Then for a centered aperture the correct formula is simply:

$$\text{angle-of-view} = 2 \arctan\left(\frac{d}{2b}\right)$$

To correctly handle an offset aperture, it is necessary to determine the offsets from image center of the exit pupil (the apparent position of the aperture as seen from the back of the lens), to specify those offsets as d and e , and also to tweak the $hfov$ slightly to compensate for the slight difference in total angle of view between the centered and offset positions. (The values of these parameters can be computed, but the formulas are non-intuitive and probably not worth the trouble to use. As a matter of practice, the Panorama Tools’ optimizer can determine proper values, given a reasonable number of control points.)

Again, dealing with offset aperture is an excellent model problem to drive and test our understanding. In practice, using an offset aperture would usually be a bad idea. The geometry model of Panorama Tools cannot separate the center of distortion from the center of projection, so it cannot deal simultaneously with both an offset aperture and non-zero $a/b/c$ corrections for lens distortion. In addition, placing an aperture off the optical axis generally selects rays that undergo worse lens aberrations, so that image quality drops compared to a centered aperture. An offset aperture might reasonably be used to create better *boke* for a mirror lens, by replacing a large donut-shaped aperture with a small circular one, but it seems unlikely that this would be used in a situation where the offset would matter to stitching accuracy.

Effects of Vignetting

Our discussion would not be complete without explaining one of the more confusing experimental results.

We noted early in this article that “it was possible to set up conditions with an external aperture such that the no-parallax point appeared to move smoothly over a distance of several inches as the built-in lens aperture was changed from fully open to fully closed.”

How could this occur? The answer is “physical vignetting”.

The experiment involved placing an 4mm external aperture on the front of a Sigma 105mm lens and progressively stopping down the lens from f/11 to f/45 while observing the relationship of foreground/background objects as the camera was rotated around

⁷ <http://www.angelfire.com/ca2/tech3d/images/persp.pdf>, pg.10.

various points. The experimental uncertainty was small, and the results were quite definite: when the built-in lens aperture was $f/11$, the no-parallax point was located at the external aperture; when the built-in lens aperture was $f/45$, the no-parallax point was located at the built-in aperture; and for built-in aperture settings of $f/16$ through $f/32$, the no-parallax point moved smoothly between those positions.⁸ These results were confusing, to say the least! How could they be explained?

Fortunately, we now understand enough to completely explain these results. On this lens, such a small external aperture produces significant vignetting when the built-in aperture is stopped down as far as $f/45$. So, for the sake of consistency, the rotation angle was chosen to place the test points just barely within the vignette at $f/45$. This left the points visible at all f /stop settings. By itself, this seems reasonable enough. But remember that the center of perspective – the no-parallax point – is established by which rays get through the lens. In placing the test points as described, we accidentally guaranteed that at $f/45$ those rays were selected by the built-in aperture, at $f/11$ they were selected by the external aperture, and at intermediate f /stops, they were selected partly by each aperture with the fraction varying smoothly. The apparent no-parallax point – the effective location of the entrance pupil – simply moved to match.

There is perhaps a useful lesson here, in addition to resolving a weird experiment.

Whenever significant physical vignetting occurs, the location of the limiting aperture changes across the image width, and so may the location of the no-parallax point. If it does, then in fact the lens will not have a single no-parallax point, but rather a collection of “least-parallax” points that vary with angle away from the optical axis, like a fisheye lens. This is not a good thing.

The practical advice is just this: to get perfect stitching, stop down your lens enough to avoid physical vignetting. You would want to anyway, to avoid intensity banding, but it’s comforting to know that you’ll be avoiding possible stitching errors at the same time.

Summary

We have explained and illustrated the following points.

The perspective of an image is determined by the light rays that formed it. Because the aperture selects those rays, its location determines the perspective. The center of perspective and the no-parallax point are one and the same, and are located at the center of the entrance pupil. The position of the optic designer’s “front nodal point” is irrelevant. The angle of view (Panorama Tools $hfov$) is measured around the center of perspective, the no-parallax point. Physical vignetting should be avoided because it may cause the lens to have only a collection of “least-parallax” points instead of a single no-parallax point.

This analysis explains all effects noted to date. Please let us know if you find a new one.

⁸ See <http://www.janrik.net/PanoPostings/NoParallaxPoint/Sigma105mmPlus4mmExtAper.gif>.

Acknowledgements

Quite simply, this article owes its very existence to the Panorama Tools forum. Special thanks are due to Erik Krause, John Houghton, David Sykes, and Alain Carrie for stimulating discussions, careful experiments, insightful suggestions and pointers, careful critiques of numerous drafts, and general good humor and open-mindedness in the face of recurring confusion (of which the author is happy to claim more than his fair share!). It was a pleasure thrashing through this problem with such fine people.

Related Reading

<http://wiki.panotools.org> The Panorama Tools wiki, an excellent entry point into the world of digital panorama photography.

<http://www.angelfire.com/ca2/tech3d/images/persp.pdf> “IMAGE-SIDE PERSPECTIVE AND STEREOCOPY”, by John Bercovitz. A mathematical treatment of many of the same ideas discussed here. The official 1998 SPIE publication can be accessed at <http://www.imaging.org/store/epub.cfm?abstrid=29536> (for money) or at <http://www-eng.lbl.gov/~bercov/optical/SPIEpaper.PDF> .

<http://physics.tamuk.edu/~suson/html/4323/thick.html> “Thick Lenses”, by Dan Suson. An introduction to the thick lens model.

http://www.winlens.de/en/lensde_opticalimg.html An illustration of the thick lens model and symbols used by the WinLens software of Linos GmbH.

http://doug.kerr.home.att.net/pumpkin/Lens_Points.pdf “Lens Principal and Nodal Points”, by Douglas A. Kerr. A readable discussion of the thick lens model.

<http://toothwalker.org/optics/vignetting.html> “Vignetting”, by Paul van Walree. An extensive discussion of various kinds of vignetting and their effects.

http://doug.kerr.home.att.net/pumpkin/Pivot_Point.pdf “The Proper Pivot Point for Panoramic Photography”, by Douglas A. Kerr. A readable discussion covering some of the same material discussed in this article.

<http://toothwalker.org/optics/misconceptions.html> “Misconceptions in Photographic Optics”, by Paul van Walree. Number 6 is a concise summary of the distinction between entrance pupil and nodal points.

<http://www.panotools.org/mailarchive/msg/25864> November 2004 discussion, linked here to provide some history of a typical debate.

http://www.photo.net/bboard/q-and-a-fetch-msg?msg_id=003bjv “Large format photography Forum: nodal point of lens”. Historical discussion, 2002.

<http://forums.dpreview.com/forums/readflat.asp?forum=1030&message=6034376> .
Historical discussion, 2003.

<http://www.luminous-landscape.com/essays/bokeh.shtml> “Understanding Boke”, by Harold M. Merklinger. A discussion of the affect of aperture shape and lens aberrations on out-of-focus images.

<http://lpc1.clpccd.cc.ca.us/lpc/molander/PDFs/Stops.pdf> Coursework presentation about limiting apertures and stops, chief rays and other concepts.

http://www.swissarmyfork.com/lens_table_1.htm A table of the position of the entrance pupil for an impressive list of lenses (unfortunately no longer being updated). Linked from http://www.swissarmyfork.com/photo_resources.html

Last edit May 11, 2009, to clarify a technical point about the thick lens model and correct a couple of obsolete URLs.