

# Ultra-High Resolution Digital Mosaics

J. Brian Caldwell, Ph.D.

## Introduction

Digital photography has become a widely accepted alternative to conventional film photography for many applications ranging from sports to wildlife to portraiture. However, two areas where most digital cameras fall short are landscape and architectural photography, where extremely high resolution, wide-angle lenses, and the ability to use camera movements to control perspective are highly desirable. Many photographers stick to using medium or large format film for these applications, and others resort to using very expensive large format digital scanning backs. Another possibility is to use a conventional digital camera to take an array of images, which are then stitched together with specialized software to form a seamless high-resolution photograph. I was first motivated to try this technique after visiting Max Lyons website:

<http://cgibin.rcn.com/maxlyons/> . Max uses a Nikon 990 to create stunning landscape and architectural images having a resolution of 20 megapixels or more. I was quite inspired by Max's work, and decided to learn the technique for myself. The amazing thing is that the sophisticated stitching software needed to accomplish this is free. It's called Panorama Tools, written by Helmut Dersch, and is available at <http://www.fh-furtwangen.de/~dersch/> . Panorama Tools also includes a very useful Photoshop plug-in that can be used to correct distortion, color fringing, perspective errors and numerous other things. Although Panorama Tools is quite powerful, it is also very hard to use. Fortunately, Joost Nieuwenhuijse has written an easy-to-use graphical interface for Panorama Tools called PTGui, which is available at <http://www.ptgui.com> .

## Taking the Images

The basic idea is to take an orderly sequence of images, each having about 30-50% overlap with its neighboring images. Having a sufficient amount of overlap is important for accurate stitching, so it doesn't pay to skimp here. The images can be taken freehand, but using a spherical panorama head like those manufactured by Kaidan ( <http://www.kaidan.com> ) will eliminate parallax error by keeping the nodal point of the lens stationary. This is not so important if everything in the scene is far away, but it becomes vital when some objects are within about 10 feet of the camera or closer. Using a good panoramic head also speeds up the taking process, which helps to ensure that the lighting conditions remain constant for all images. In the example I'll be discussing I used a Nikon D1 camera to take a mosaic of 10 images in a pattern of two rows of 5 images. Figure 1 shows what these images look like if they are simply laid out in an array. I used a manual focus 35mm f/1.4 lens stopped down to f/8, and the exposure for each image was held constant at 1/800 sec. It is important to have an identical exposure for all the images in order to avoid visible seams in the final mosaic. All 10 images were taken in a span of only 25 seconds. Framing the photograph involves quite a bit of

guesswork. I used a piece of cardboard with a rectangular cutout to frame the scene, and then setup the shoot to make sure that the images would fill-out the scene that I had framed.

I try to take RAW images with the D1 as often as possible, and process these images in Bibble. This provides excellent color and maximum exposure flexibility. For relatively small mosaics like this one, RAW is a viable option because the camera buffer doesn't hinder the shooting. However, for bigger mosaics of 20 or 30 images it is sometimes necessary to use JPEG-fine in order to be able to take the images as rapidly as possible. When shooting 8-bit JPEGs it is a good idea to take test shots in several areas to be covered by the mosaic to be sure that the exposure is optimal.




Figure 1: Array of D1 images used for constructing the mosaic

### **Preparing the Images**

The 35mm lens that I used is fairly high quality, but it does have some distortion. I had previously calibrated this lens by photographing a doorway so that one of the straight edges was very close to the side of the image frame. I then varied the “b” and “d” parameters in the Radial Shift correction option of the Panorama Tools plug-in to make the door edge as straight as possible. After taking the mosaic images I then batch-corrected them using these parameter values, as shown in Figure 2.

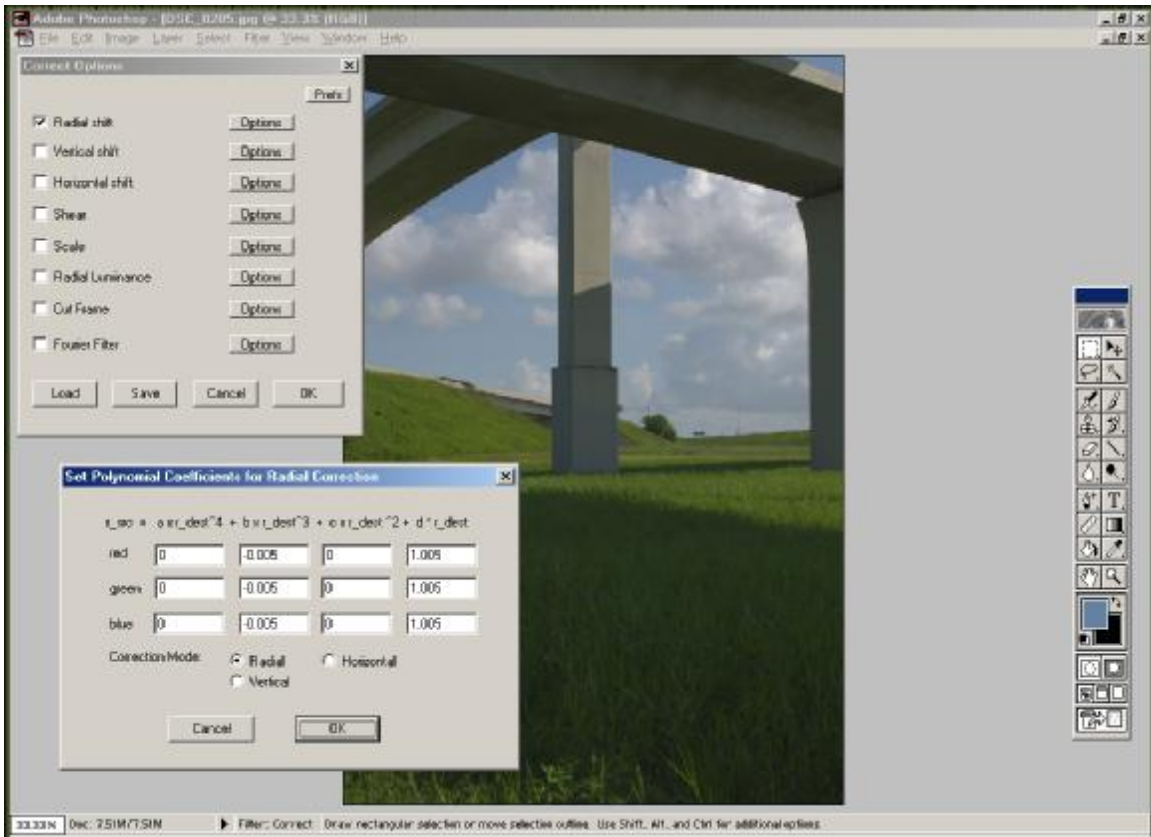


Figure 2: Correcting the distortion of one of the images

### Stitching the Mosaic in PTGui

After eliminating all traces of distortion using the Panorama Tools plug-in, the images are added to a new project in PTGui as shown in Figure 3.

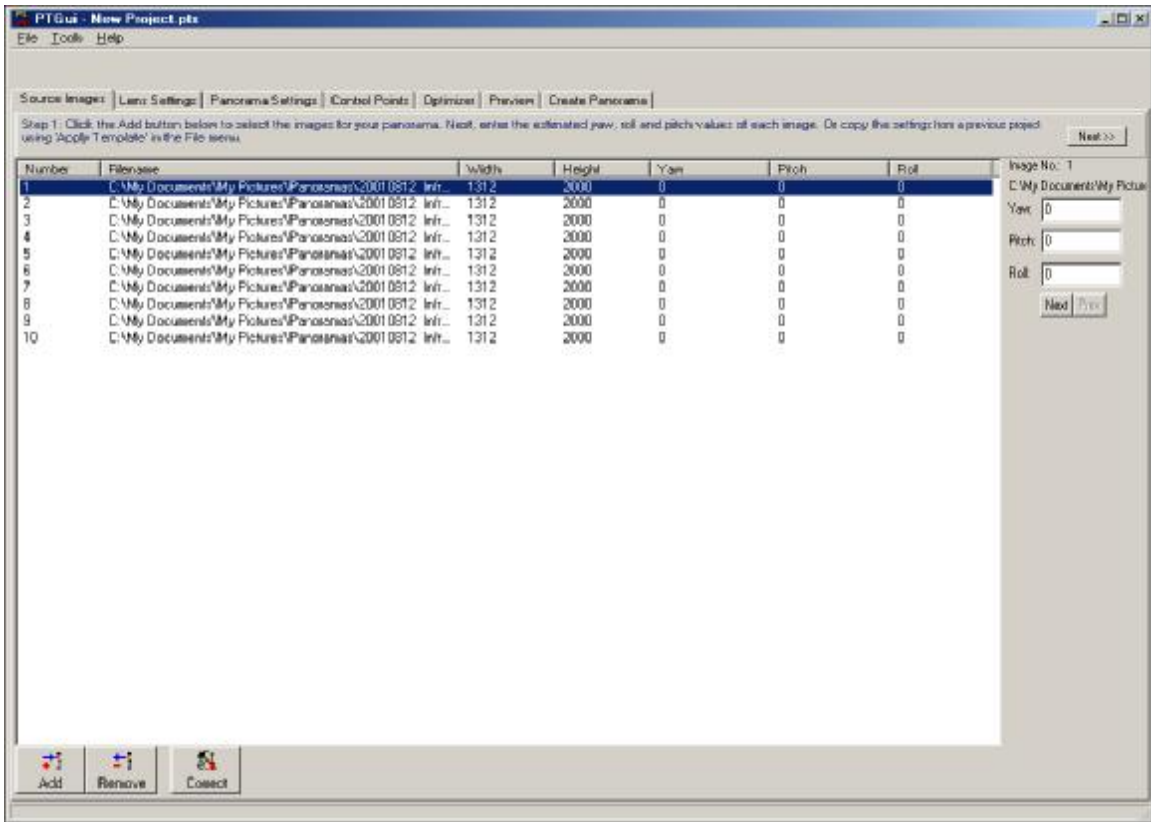


Figure 3: Adding the images to PTGui

Next, the initial lens parameters are entered as shown in Figure 4. The horizontal field of view is calculated along the short side of the image frame since the images were taken in portrait orientation. The D1 imager is approximately 15.5mm x 23.6mm, so for a 35mm lens the horizontal field of view will be  $2 * (\arctan((15.5/2)/35)) = 25$  degrees. This value will be optimized later on to account for uncertainty in the actual focal length of the lens. The distortion parameters a, b, and c will be held fixed at zero because distortion has already been removed from the images.

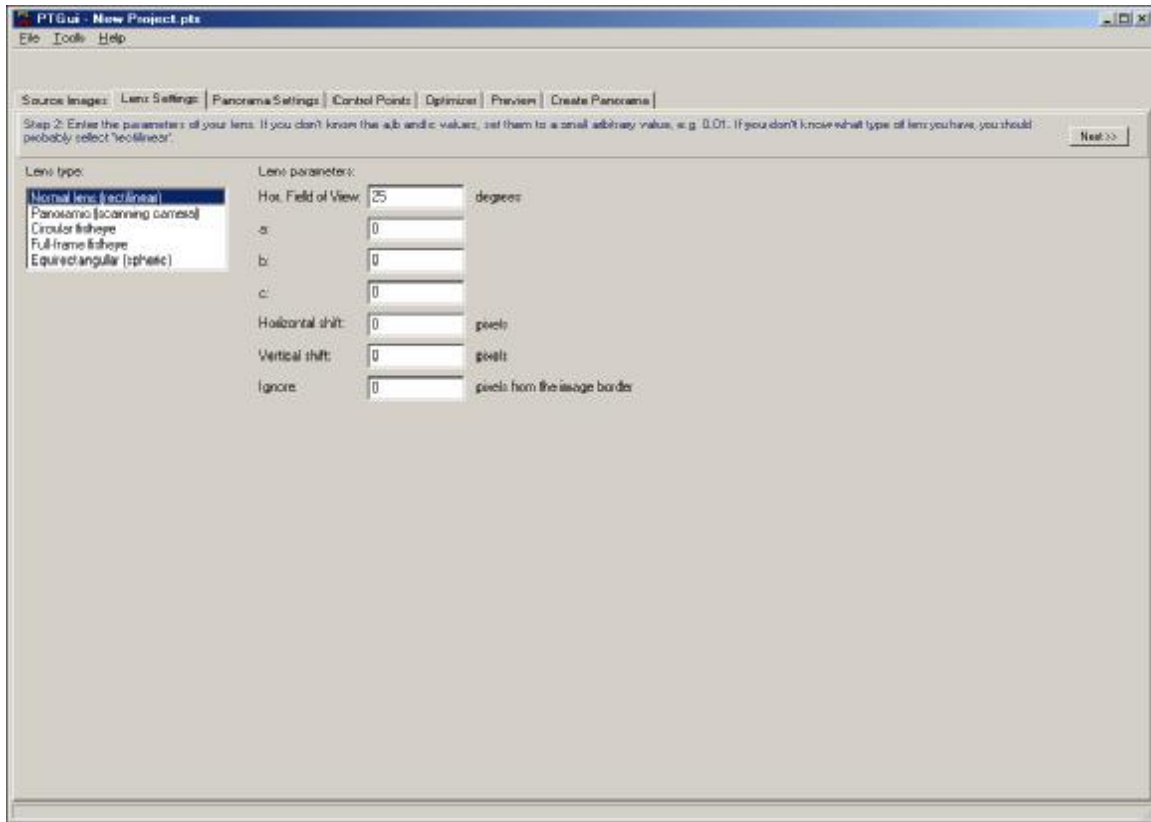


Figure 4: Lens settings

The initial panorama settings are now entered. The width, height, and field of view will be changed later after checking some preview mosaics.

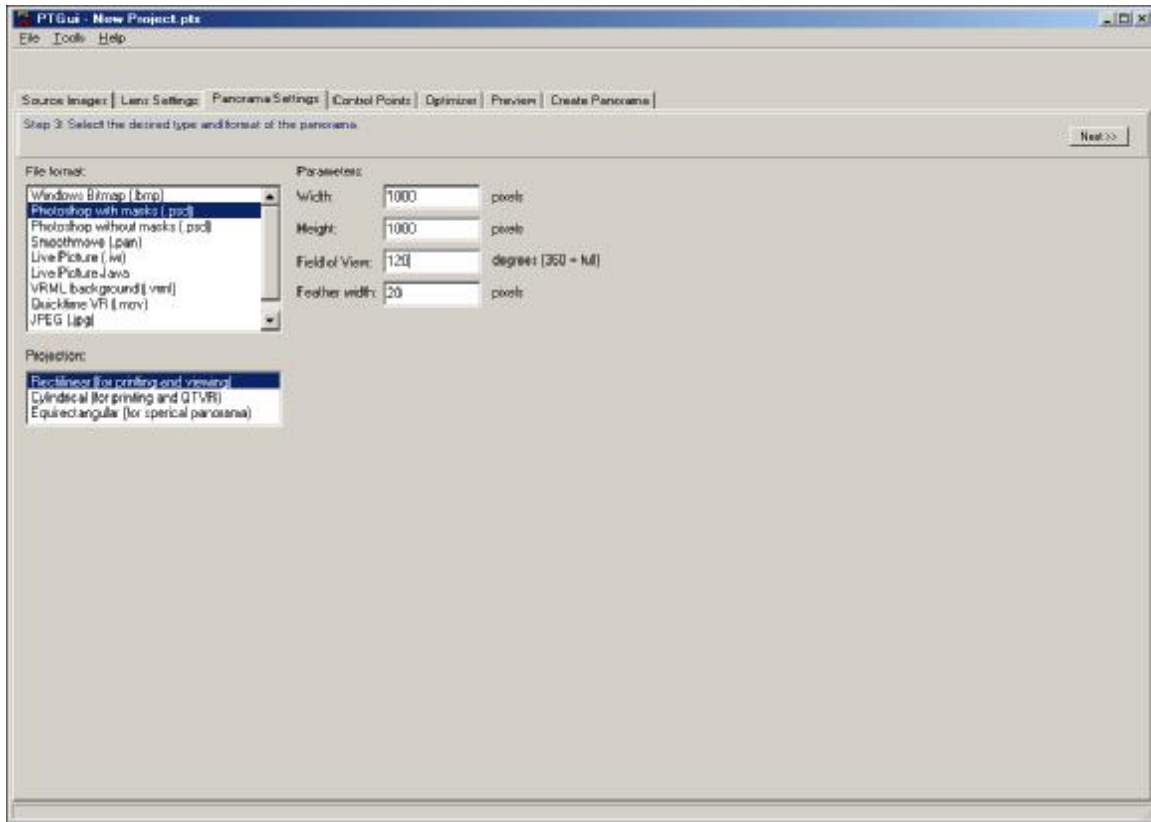


Figure 5: Initial panorama settings

In order to get an accurate stitch it is necessary to manually identify common scene elements in each pair of neighboring images. These scene elements are called control points, and they must be identified between adjacent images in different rows as well as between adjacent images within the same row. For a 2x10 mosaic there are 13 pairs of images for which control points must be established. I try to find a minimum of 4 widely spaced control point pairs in each overlapping pair of images. Figure 6 shows the PTGui window used for finding and marking the control points. The two large windows show a small area of the overlap region between two adjacent images. The small window in the bottom right is a magnified view that helps in achieving single pixel or better accuracy in setting the control points.

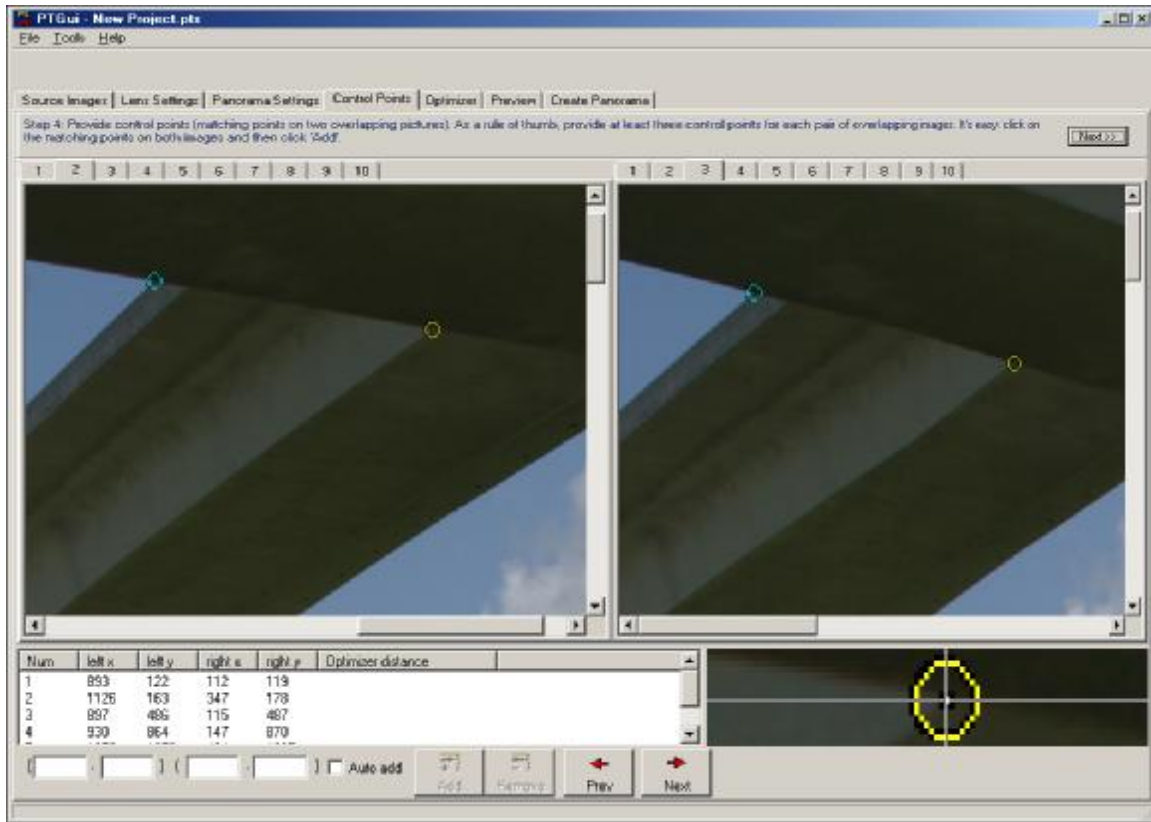


Figure 6: Setting control points between each facing pair of images

Next, the pitch, yaw, and roll angles of the images must be optimized to minimize the average distance between the control points. Yaw is the azimuthal, or left-to-right angle of the camera. Pitch is the elevation, or up-and-down angle of the camera. Roll is the rotation angle about the lens axis. I find that it is best to use the Panorama Tools optimizer in stages to get best results. Figure 7 shows how the first optimization run was set up. First, I allow all of the yaw angles to vary except that of image #3, which will be used as a reference image. Since the first row of images was taken with the tripod head in a level position the pitch angle of images 1-5 will be left at zero while the angles of the second row of images will be allowed to vary. The roll angles will be initially left at zero because the camera was set up perpendicular and wasn't allowed to rotate about the lens axis. This initial optimization run gave an average distance between control points of only 0.5 pixels, which is reassuringly small. Figure 8 shows a 1000 pixel wide by 850 pixel high low-resolution preview of the mosaic after the first optimization run.

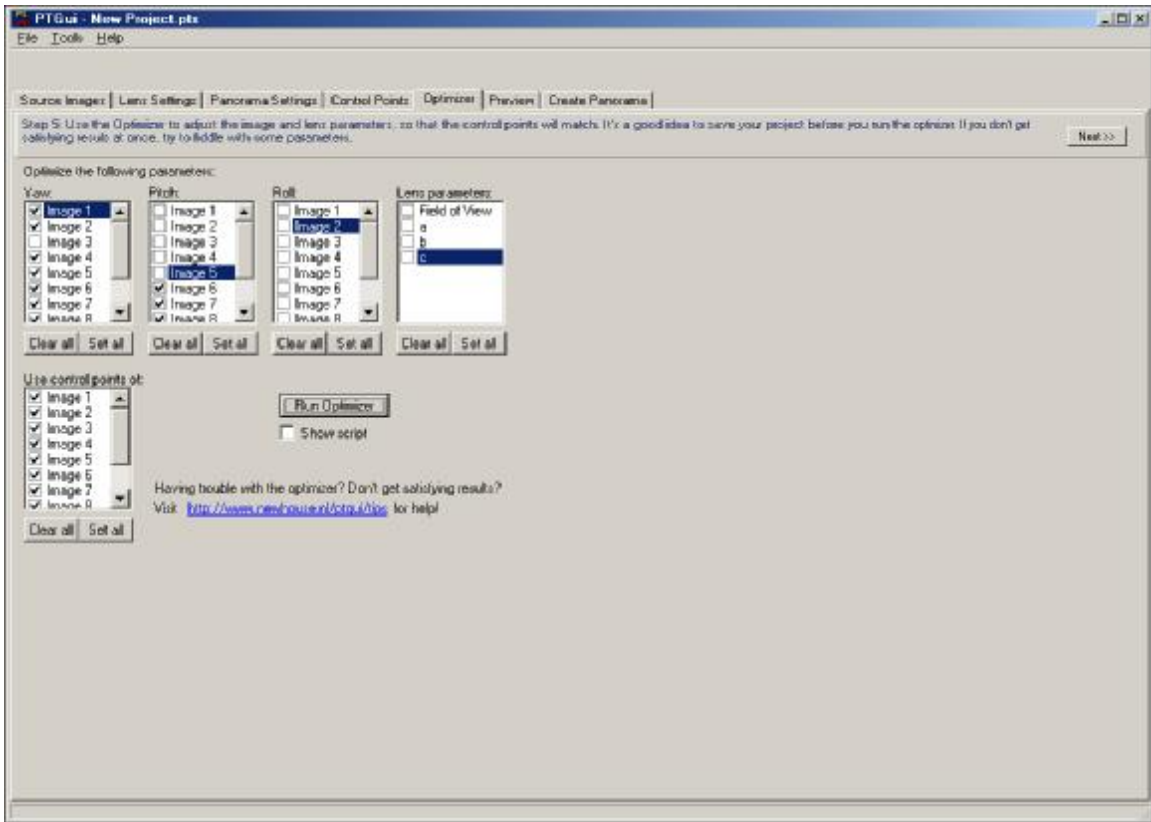


Figure 7: First optimization run – gives an average distance between control points of 0.5 pixels





Figure 8: Preview mosaic from the first optimization

One reason for having a reference image is that you can manually change its orientation to change the image perspective. In the Figure 8 preview I didn't care for the exaggerated size of the roadway in the upper right relative to the one in the upper left, so I decided to even things out a bit by rotating the entire image to the left. I accomplished this by fixing the yaw angle of the reference image to  $-10$  degrees as shown in Figure 9, and then re-optimizing using the parameters set in Figure 10. Note that here I've allowed all of the angles to vary except those of the reference image, and in addition I've allowed the field of view to vary. This second optimization reduced the average distance between control points to an insignificant 0.1 pixels, and a preview is shown in Figure 11.

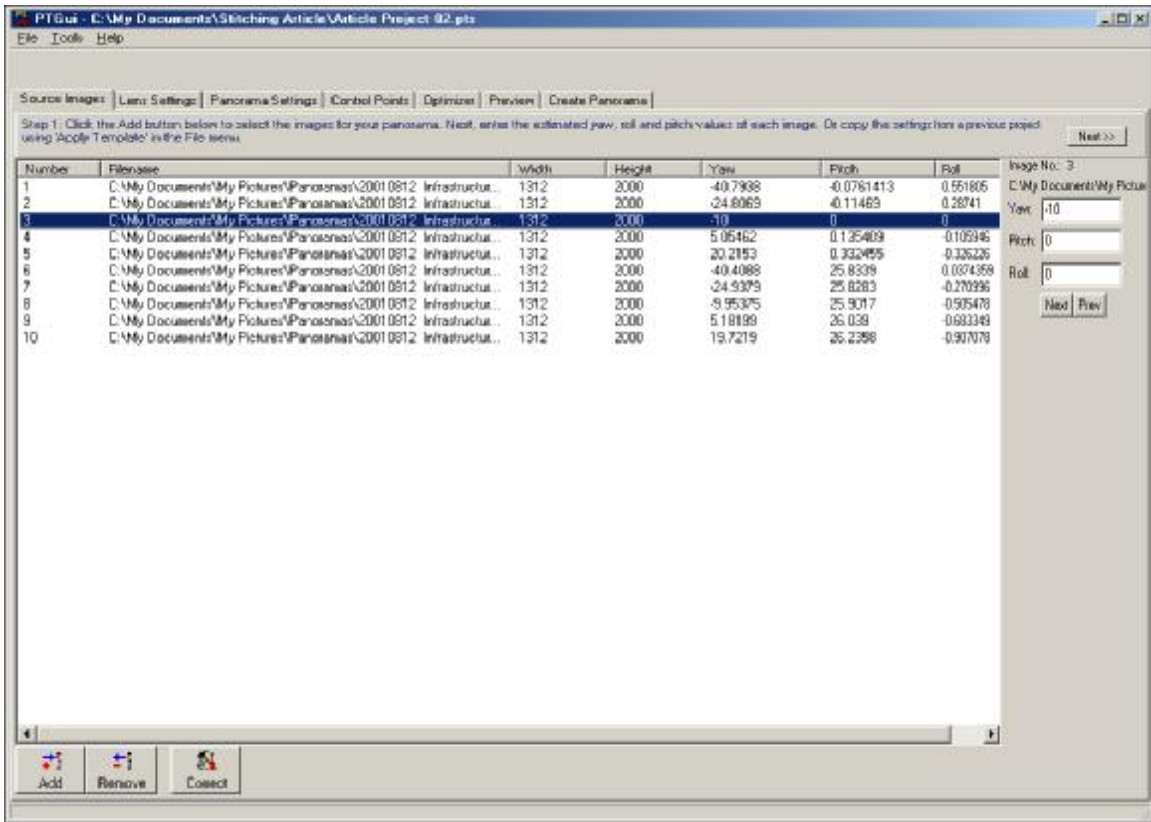


Figure 9: Manually adjusting the yaw (left-right) angle of the reference image

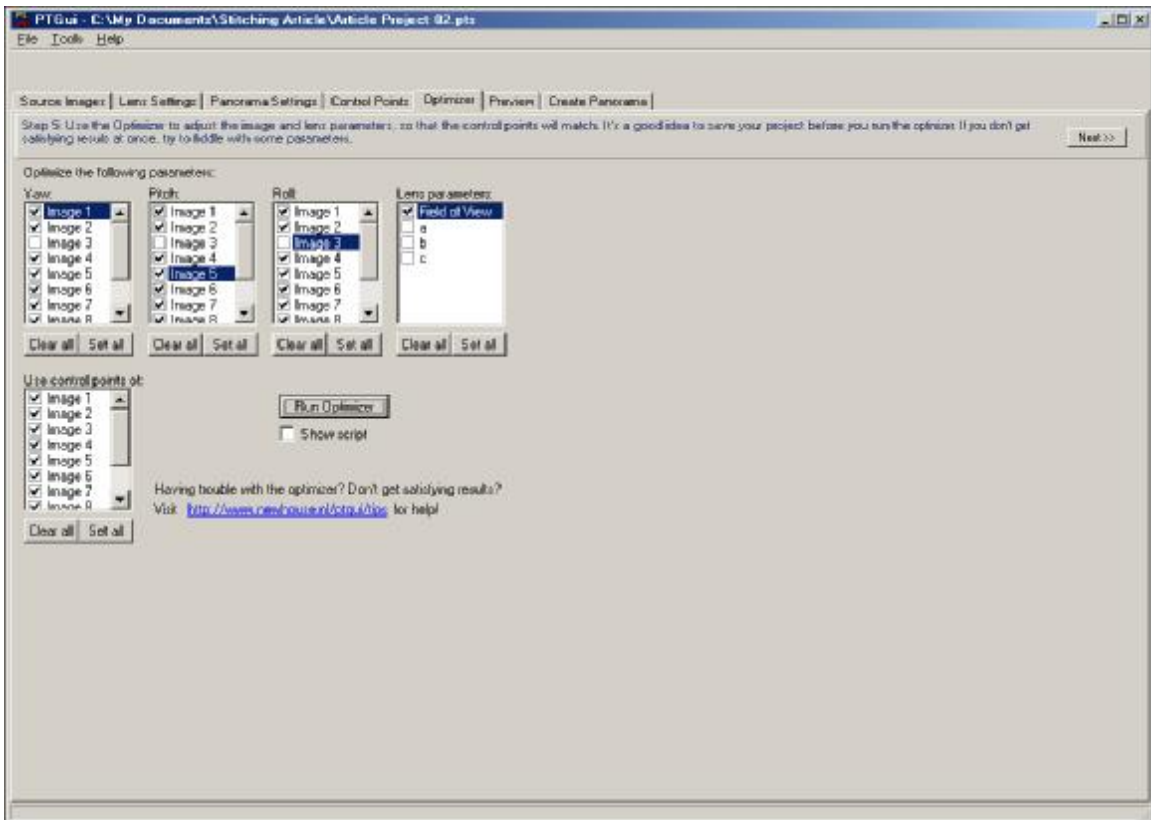


Figure 10: Second optimization run




Figure 11: Preview after second optimization

The large black-colored area around the image in Figure 11 means that the 120 degrees that I initially chose as the horizontal field of view for the entire mosaic is too large. So I reduced this to 105 degrees as shown in Figures 12 and 13.

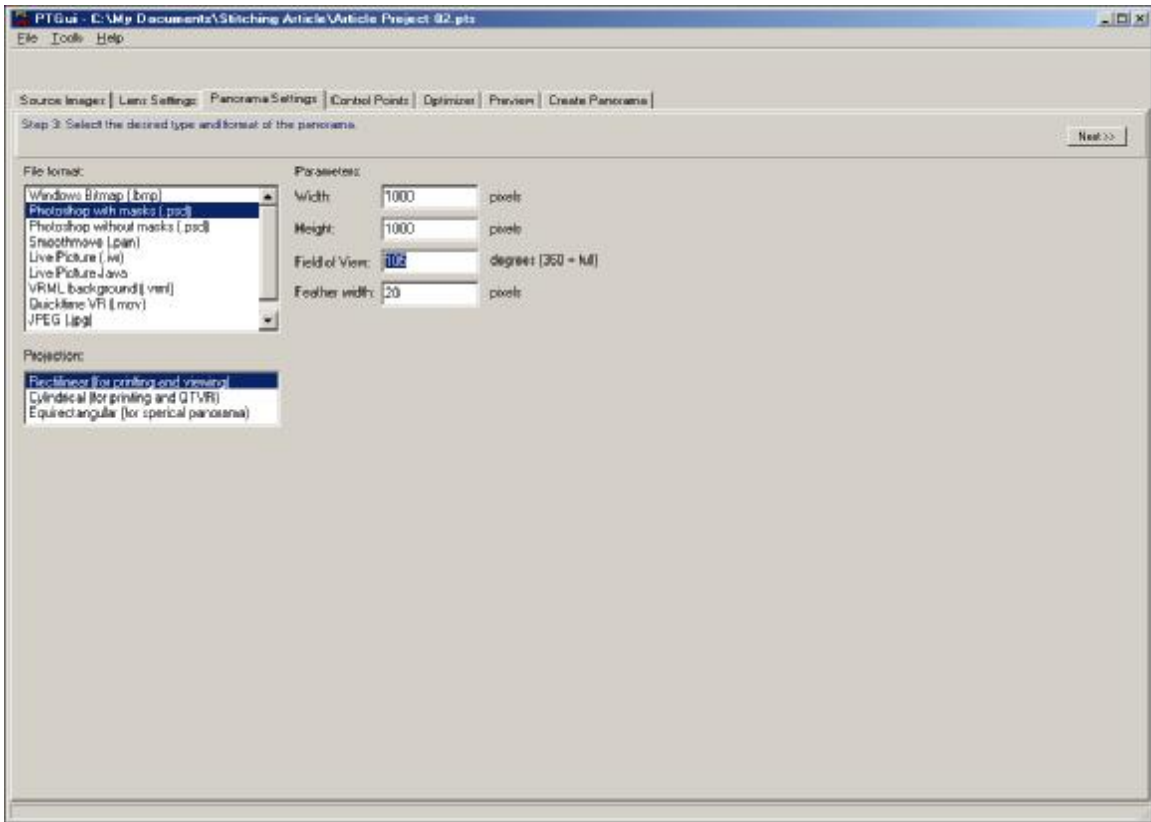


Figure 12: Reduce the field of view to eliminate some of the blank area surrounding the image



Figure 13: Preview with reduced blank area

The final step before calculating the full size mosaic is to determine what the optimum resolution should be. To do this I go to Photoshop and open up the preview that I am satisfied with along with one of the original images. I prefer to choose an image area near the middle where the stitching process doesn't alter the geometry too much. In Figure 14 I have had to enlarge the preview to 200% in order to match the visual size of scene elements in the original image at 25%. Since the preview is 1000 pixels wide by 850 pixels high, I will need to change the final panorama width to  $1000 * 200 / 25 = 8000$ , and the final height to  $850 * 200 / 25 = 6800$ , as shown in Figure 15.




Figure 14: Comparing one of the original images to the mosaic preview. The preview must be enlarged about 8 times more than the original image to make scene elements appear the same size.

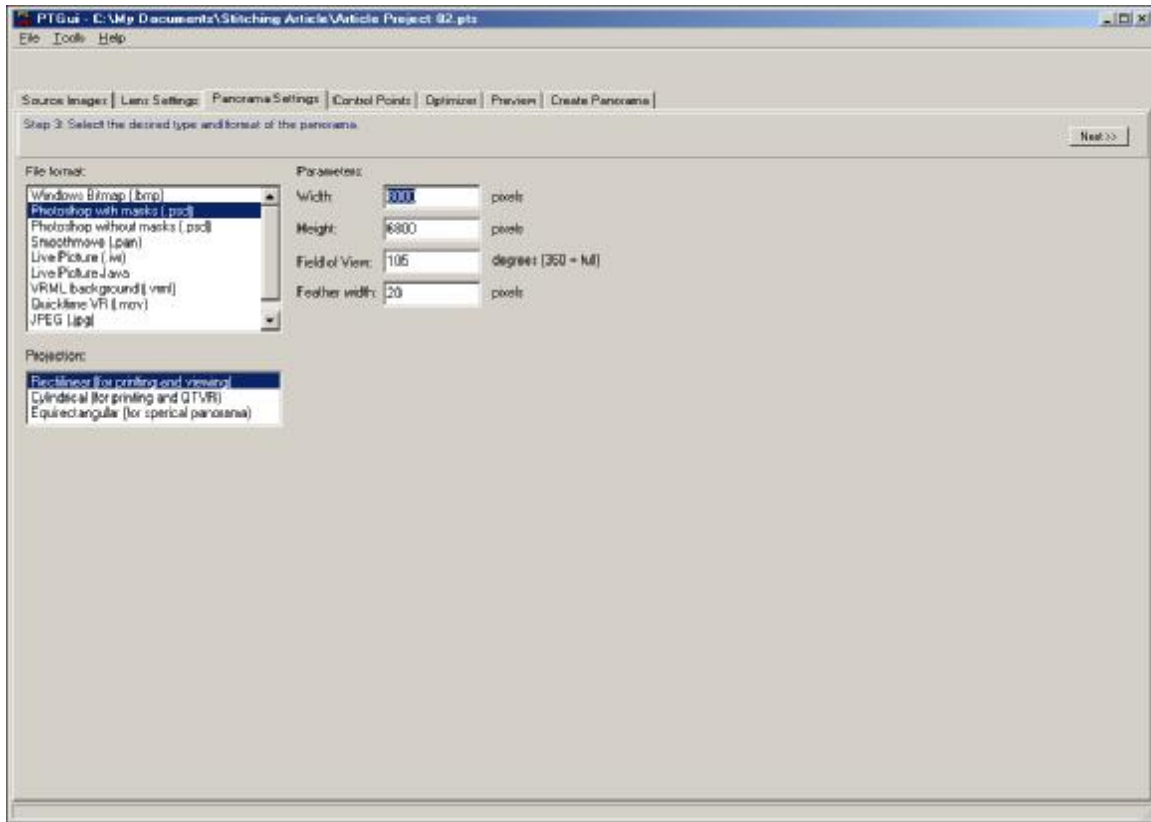


Figure 15: Setting the height and width of the panorama

The next step is to actually stitch the image together and save the result to a file, all of which is accomplished without any manual intervention. This step took about 45 minutes to complete using a 1Ghz Pentium III machine with 512MB of RAM. The output file in this case is a multi-layer Photoshop (.psd) file of just over 200MB. The advantage of having the layers is that any visible seams caused by drifting subject matter, slightly varying lighting conditions, or changing amounts of flare due to pointing the lens in different directions can be eliminated by airbrushing the layer masks. In this case I only had a small amount of touch-up to do, all of which was accomplished in a few seconds by airbrushing the layer mask of layer #10, as shown in Figure 16.




Figure 16: Eliminating visible seams by airbrushing one of the layer masks in the multi-layer Photoshop file output by Panorama Tools.

Next, I corrected a small amount of residual perspective and rotation error using the very intuitive LensDoc plug-in published by Andromeda, as shown in Figure 17.





Figure 17: Correcting slight perspective and rotation errors using the Andromeda LensDoc plug-in.

The final mosaic after cropping and slight curves manipulation is shown in Figure 18, and an enlarged detail is shown in Figure 19. This mosaic is 5937x3969, or 23.6 megapixels, and it is similar in quality to what might be achieved by using a 35mm ultra-wide angle shift lens on 6x7 format. To get a similar perspective on a single D1 frame would require something like a 12mm shift lens! One interesting benefit of the mosaic technique over conventional ultra-wide angle photography is that there is virtually no illumination falloff in the corners.



Figure 18: Final seamless mosaic after cropping and curves adjustment




Figure 19: Detail from the pink-shaded area in the final mosaic

### Further Reading Online

Helmut Dersch's site, <http://www.fh-furtwangen.de/~dersch/>, has many interesting articles and links, and is well worth a detailed look. An excellent tutorial by John Houghton on using Panorama Tools to create printable panoramas is located at <http://homepage.dtn.ntl.com/j.houghton/pttute.htm>. <http://www.panoguide.com> is an excellent source of information on panoramic photography in general and stitching software in particular. Another excellent panorama site and tutorial is <http://www.bigbenpublishing.com.au/360/docs/tutorial/advanced/index.html>.