Hardware & Software Considerations for the Creation of HDR Panoramas

Dr Ryan Southall - School of Architecture & Design, University of Brighton.

INTRODUCTION

High Dynamic Range (HDR) imaging is a digital photographic technique that composites together a number of differently exposed conventional digital images to produce an image with a higher and more accurate brightness range, which is numerically coded into the image format. Common HDR image formats are Radiance RGBe (.pic/.hdr extension) and OpenEXR (.exr extension). Formats differ in the coding of the brightness data, and the numerical range available for this coding. This photographic technique was first revealed in 1997 [1] and initially intended for the generation of realistic lighting sources for virtual environments. The use of HDR imaging has since become popular in the general photographic world as it allows somewhat surreal, or 'other worldly', images to be created. For either purpose the generation of HDR images is essentially the same. This document details hardware and software choices that can cost-effectively optimise the capture and generation of HDR images.

GENERATING HDR IMAGES

Auto-Exposure Bracketing

To generate an HDR image a number of differently exposed digital images of the same scene must be created. This task has been greatly simplified in recent years with the addition of an Auto Exposure Bracketing (AEB) feature to digital SLR's and high end compact cameras. AEB automatically takes a number of user specified shots with a user specified Exposure Value (EV) difference between them. This variation in exposure should be achieved by variation in shutter speed, not aperture, and the camera is therefore usually in Aperture Priority, or AV, mode when the images are taken. This results in a number of differently exposed, bracketed shots (Figure 1). Key features of a camera for HDR photography are therefore the number of bracketed shots that can be taken, and the EV range between them. Three, five (for lower end cameras), seven and nine (higher end) are typical values for number of shots with an EV step between each shot of between 1/3 to 2. The higher either value the greater the brightness range that can be covered, and in general the greater cost of the hardware. A Samsung gx-20 camera is used on this project as it can take 5 bracketed shots with a 2EV step between each shot, giving a total range of 8EV. The actual brightness range in a sunlit scene may equate to up to 20EV steps [2], and it is unlikely that commodity hardware will be able to capture this whole range in the near future, but the higher the total EV range that can be covered the better, as more of the brightness range of the scene can be captured. If the scene is an interior, or overcast, scene with no strong sources of illumination then the number of bracketed shots can be reduced and the shutter speed times centred around the normal exposure level. If there is direct sunshine in the scene the camera's shot bracket number, EV step and aperture value should be increased to their maximum values (i.e. the aperture is as small as possible) and the exposure level set to the minimum setting (i.e. the shutter times are as short as possible).

Image resolution is a decision for the photographer with high resolution images delivering a higher resolution final panorama but requiring more computer memory and time in the HDR creation phase.

As a number of identical positional shots of the scene are taken within a bracketed set it is advantageous if the camera does not move between shots, as any misalignment may have to be corrected in the post-processing phase. Movement of the camera, and subsequent misalignment, can often occur during the first shot as the camera button is pressed, so using a remote control to avoid touching the camera, whilst shooting image sequences, improves the alignment of the images.

Figure 1: Sequence of bracketed shots.

The next hardware consideration is the camera lens, and this choice is dependent on the chosen panorama photographic technique. To wrap the generated HDR image around 3D virtual space the image must capture the lighting from the whole 360^o of the scene, and generating these 360^o panoramas can be achieved in one of the two ways: mirror ball and panorama stitching techniques.

Mirror Ball Technique

Shooting an image of a mirrored sphere can theoretically capture lighting data from the whole scene, including the reflection of the camera and photographer (Figure 2).

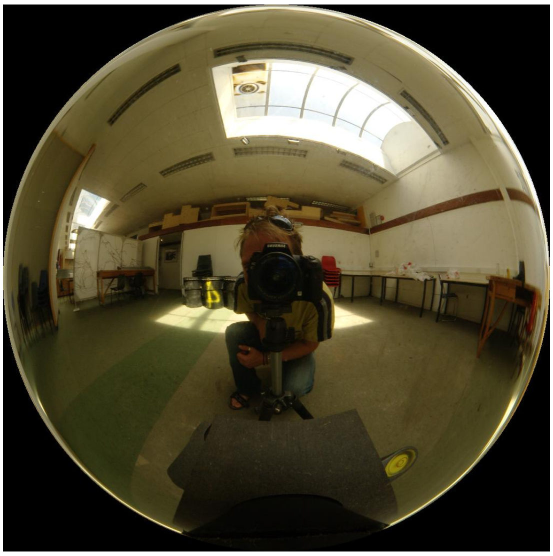

Figure 2: Mirror ball shot.

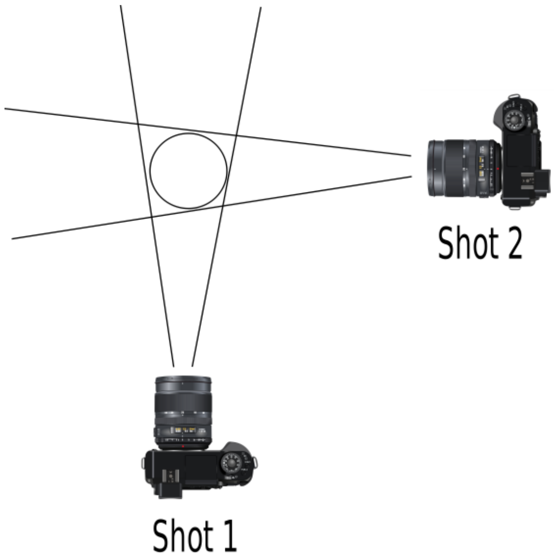

Figure 3: Shooting mirror ball to remove photographer.

However, as a higher angular range of the scene data has to be represented by a smaller visible surface of the sphere as the view at the edge of the sphere is approached, the resulting data near the periphery can be distorted. This distortion can be reduced, and the image of the photographer and camera removed, by taking two sets of bracketed shots of the sphere at a 90° angle from each other (Figure 3). One of the resultant HDR images is then mathematically rotated by 90° and can then be used to fill in the distorted area, or reflection of the photographer, in the other.

This was the technique initially investigated for this educational context as it has a number of advantages.

- Cost – effective: Hollow chrome steel spheres or bearings cost in the region of €15.

- No special lens is required.
- Only one or two sets of shots are required.

There are however a number of disadvantages of the technique including:

- Two tripods are required
- Imperfections in the sphere effect image quality
- Difficulty in achieving accurate 90° projection transformation
- Not suitable for uneven terrain
- Relatively complex software process is required to produce a final panorama.

After testing, the disadvantages of the above technique were significant enough to warrant investigation of the second technique, panorama stitching.

Panorama Stitching Technique

The stitching panoramic technique requires conventional images of the whole scene to be taken. This can be done with many images and a conventional lens, or fewer images and a wide-angle or fish-eye lens. The advantage of a fish-eye lens, and taking fewer images, is that as the image acquisition process is quicker lighting levels are less likely to change, due to cloud cover etc and fewer images need to be stitched together requiring which requires reduced computing resources. On this project a Samyang 8mm aspherical fish-eye lens has been used, which can be bought for around £200. Other makes include Rokinon and Bower. An aspherical 8mm lens takes a rectangular image with a 180^o view angle from diagonal corner to corner. In portrait mode the horizontal view angle is 100^o and the vertical is 150^o. The 100^o horizontal view angle means that 4 sets of bracketed portrait shots, with a 90^o horizontal angle between each, can be taken and just cover the 360^o horizontal angular range of the scene (Figure 4). Care must be taken however during shooting that the images are taken in increments of close to 90^o; this may be achieved by orientating the camera between shots relative to a landmark in the scene or aligning with the sides of a rectilinear shape, such as a sheet of paper placed under the tripod.

To make sure the 150^o vertical view covers the zenith of the scene the camera should be tilted up by 20^o. This does mean that a 30^o view angle below the camera is not captured but as this portion of the scene is usually obscured by the virtual model to be illuminated, in general contains little illumination and also contains the tripod, this is considered to be an acceptable trade off for sped and simplicity of image capture. Although fisheye lenses can be more costly than conventional lenses, and more images are required than with using a mirror ball, this technique has a key advantage; the resultant images are of higher quality, with less post processing effort, than the mirror ball technique. This is largely because of the popularity of panoramas within the general photography world which has led to many imaging applications now having the capability to seamlessly and automatically stitch together a number of overlapping shots to create 360° panoramas e.g. Photoshop (http://www.photoshop.com/), and, importantly, deal natively with HDR images e.g. Hugin (http://hugin.sourceforge.net/).

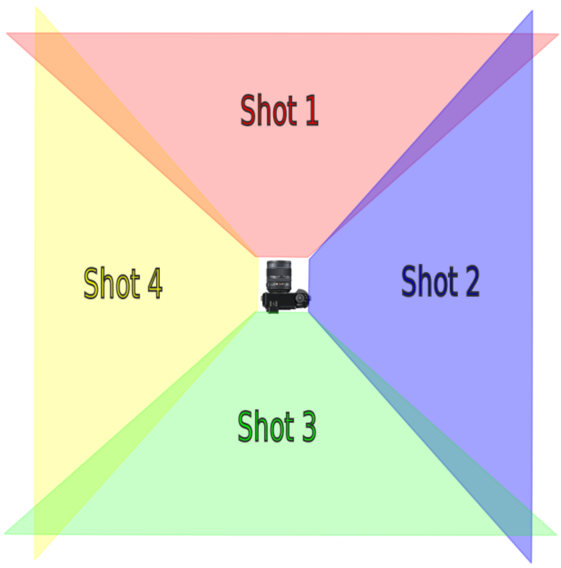

Figure 4: Panorama shooting on the horizontal plane.

Parallax Errors

If the camera is connected in a conventional manner to the tripod then as the camera is rotated the position of the end of the lens, the point at which the light in the scene is gathered, will change (Figure 5a). This movement of the end of the lens relative to the scene causes parallax errors whereby objects in the foreground are seen to move relative to the background. This creates errors in the photo stitching phase for scenes with foreground objects. To reduce or eliminate these errors the end of the camera lens should be positioned directly above the tripod pivot point. This can be achieved with a separating plate that sits between the camera and the tripod (Figure 5b). The separator plate has an embedded hex nut with a 1/4 inch standard photographic Whitworth thread at one end, and a hole to take a corresponding bolt on the other. The bolt screws into the bottom of the camera and the hex nut accepts the bolt from the tripod mount plate.

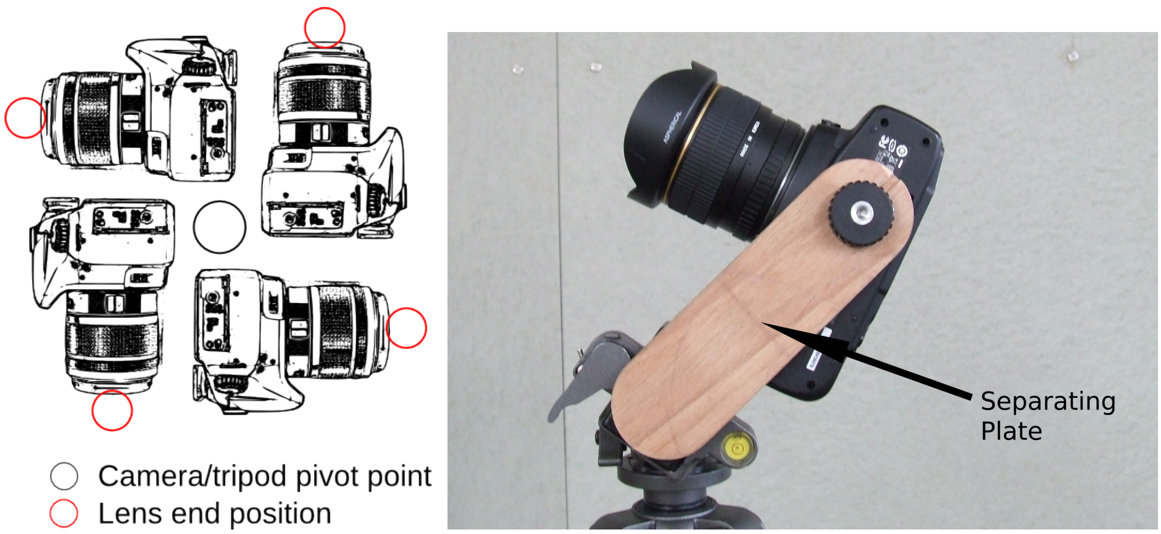

(a) Lens position change with shot rotation.

(b) Camera position correction with separator plate.

Figure 5: Cause and correction of parallax errors.

An svg vector file of the separator plate can be found in the 'Resources' section of this website.

Software Composition

Each of these conventional, low dynamic range (LDR), bracketed shots are limited in the brightness range they can represent; both by the camera sensor and the image format the camera exports, and must be composited together.

The LDR nature of conventional images has traditionally not been of huge importance due to the LDR nature of common viewing mechanisms e.g. photographic paper or computer monitors that cannot display a higher brightness range even if it was present in the image. 8-bit JPEG images that can be exported by almost all cameras only have a maximum 1:256 ratio between the darkest the brightest elements in the image. Each of these 8 bit bracketed shots therefore only contain an accurate brightness representation over a sub-set of the overall scene brightness range but because they are differently exposed, and different parts of the scene fall within the limited brightness sensitivity of the sensor in each shot. By com-

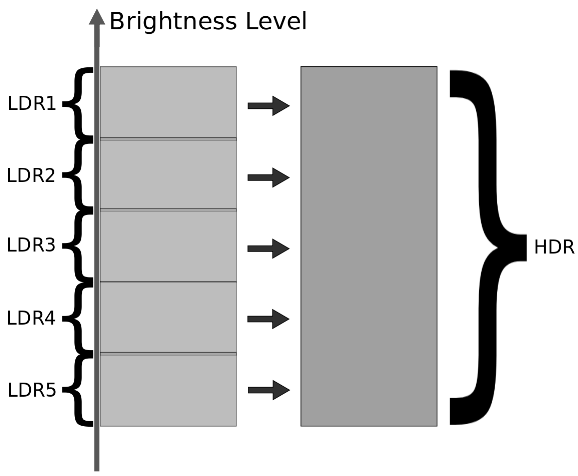

Figure 6: Combining LDR shots to increase dynamic range.

positing these images together the subsets of brightness range can be overlapped to provide a higher brightness range (Figure 6). This higher brightness range then resides in the resultant, single digital HDR image.

There are a number of software programs that can composite bracketed LDR images into an HDR including PhotoShop, Luminance-HDR, HDRShop (http://gl.ict.usc.edu/HDRShop/) and PhotoSphere (http://www.anyhere.com/) but few can also stitch a number of bracketed sets into a 360^o HDR panorama. Hugin is one of the latter, has the added advantages of being free, open-source and multi-platform and has been chosen as the software tool for this project. A more detailed description of the use of Hugin for the creation of HDR panoramas can be found in the Documentation section of the HIDE website, in both text and video tutorial form.

Tone-mapping

For reference, tone-mapping is a software-based, mathematical treatment of the numerical brightness data within the HDR image to compress the range for display on LDR systems.

Different algorithms exist to compress the range: some try to mimic the human eye response (which also has a limited brightness range), some bring out detail on the darker or lighter parts of the image and some simply scale the range so that detail in both the lighter and darker areas of the image can be viewed. It is these tone-mapping techniques that have largely led to the popularity of HDR as a general photographic technique as it allows detail in images to be seen that one does not expect to see. Detail within areas of heavy shadow can be seen next to detail within areas of high luminance whereas the human eye, or a conventional photograph, would overexpose, or underexpose, one or the other. This is what often gives tonemapped HDR images their surreal quality (Figure 7). For this project unprocessed HDR images are being used for the illumination of virtual space, but the free software tool Luminance-hdr (http://qtpfsgui.sourceforge.net/) has a good range of tone-mapping algorithms if tone-mapping is required.

Figure 7: Tone-mapped HDR image with the Mantiuk 2008 algorithm.

Final Set-up and Conclusions

The key feature desirable in a camera for HDR image creation is Auto Exposure Bracketing (AEB) in Aperture Priority (AV) mode. The number of, and EV step size between, the AEB shots is also of importance with a higher total EV range capable of capturing a wider brightness range in high contrast scenes. A total EV range of 8 is achievable with the camera used on this project, which was considered to be a good compromise between cost and performance. It is also advantageous to have a camera with remote control capability to avoid camera movement when taking shots.

The panorama photo-stitching technique offers better quality output and ease of software post-processing, which in turn makes the use of a wide-angle or fisheye lens advantageous as fewer shots are required to capture the whole scene. The exact nature of the lens is flexible but this project as used a 8mm full-frame fisheye with a 100^o horizontal view angle in portrait mode, allowing only 4 sets of shots to capture the whole 360^o horizontal scene. If a smaller view angle lens is used more shot sets may be required.

On the software side the photo-stitching package Hugin has been used for its ability to deal natively with HDR images and its free, open-source and multi-platform nature. The Hugin package is one of the few software programs that can generate HDR images (EXR format) and stitch them together into a panoramic or light probe format. Hugin can export both equirectangular projected images, which equates to the polar, or lat-long, format (Figure 8a), or full frame fisheye images, which equate to an angular light probe format (Figure 8b).

(a) Equi-rectangular, or lat-long, format HDR image. (b) Angular light probe format.

Figure 8: Light probe image formats.

Most image stitching packages work on the principle of finding common points in the overlapping portions of the images and using them to orientate the images with respect to each other, and this is also the case with Hugin. Another advantage of Hugin is that once the lens and camera combination has been calibrated, every subsequent stitching only requires a few control points. In an educational context where a student might only be producing a one-off HDR light probe, this simplification of the process is very advantageous.

After investigation into different hardware and software options the following set-up was chosen:

- Camera – Samsung gx-20 (up to 5 bracketed shots in 2EV steps for an 8EV total range)
- Lens - Samyang 8mm, aspherical, fish-eye lens
- Software - Hugin, open source, free and multi-platform panorama stitcher.

The overall cost of the system was approximately £600, which although significant, should not be prohibitive for most educational institutions.

References

[1] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. *SIGGRAPH 97*, August 1997.

[2] I. Valiev, A. Voloboy, and V. Galaktionov. Improved model of ibl sunlight simulation. In *Proceedings of the 24th Spring Conference on Computer Graphics*, pages 27–32. ACM, 2008.