

# Make/Fake an HDRI for IBL in Bryce for Bryce

## Part 2: Rendering the “Exposure” series

Now we have the scene and can render a panorama. In photography, we need a series of exposures with different settings that control the light reaching the film or sensor. For Bryce we need a series of renders under different light. We have to establish what lights we have, what intensities those lights have and also consider the 2D faces that emulate the light sources.


Once we know what we have, we can determine how we have to change the values between the renders and how many renders we will need. Each scene will have different lighting and you will have to create a table for each setup. This table is only valid for this very scene.

f-stop	Diffuse			Color Luminance	Amb/Diff RGB	Render
	A Spotlights	B Spotlights	Filler Radials	Filler Radials	2D Faces	
f/1	999	500	4	255	255/255/255	Yes
f/2	500	250	2	255	255/255/255	No
f/4	250	125	1	255	192/192/192	Yes
f/8	125	64	1	128	128/128/128	No
f/16	64	32	1	64	64/64/64	Yes

Bryce does not need the f-stop column, but programs that merge exposures to an HDRI do. f-1 is an ideal camera and does not exist in reality, but the programs accept it and it is easiest to start with this value. f-2 has the aperture one full f-stop smaller, that is, the sensor gets half the light of f-1. f-4 gives half the light of f-2 and a fourth of f-1. Equally, f-8 results in an eighth and f-16 in a sixteenth of the light if f-1.

There are 2 groups of spotlights in the scene, denoted A and B in the table. Their Diffuse value is halved for each full f-stop down. Then, there are two filler radials with no shadow casting and no falloff. Their diffuse value is also halved but beyond f-4, we cannot half diffuse again because only integer values can be entered. Therefore, we set the colour from white to mid-grey, which also halves the light, and then again to dark-grey 64, which halves the light yet again.

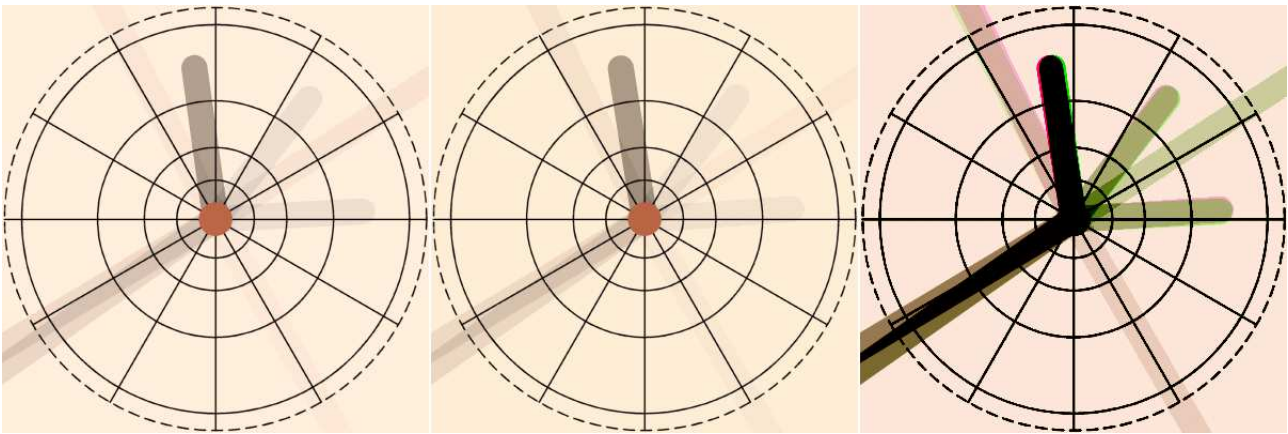
The white 2D-Faces are what is interpreted as the brightness of the light source. We can change the white for Ambient and Diffuse to some grey shades, as we did for the radials. But here we do not follow the pattern of halving brightness by making the white progressively grey because if we did (blue curve), they would not stand out as brighter than the rest (red curve).



We want them as light sources in the HDRI. The whiteness stays the same for f-1 and f-2, which will be interpreted as over exposed parts. Only at f-4 the white gets a bit grey, but only a fourth from 255 to 192, not half. This is the method we use to make the “light sources” stand out.

Finally, we decide which settings should be rendered. If we save the renders as 24 bit BMP, we would need all five renders. However, we will export the rendered images as 48 bit TIFF and the brightness resolution is high enough to use “exposures” with two full f-stops, resulting in a render time saving of 40%.


For this scene, I did render all settings, exported them as 48 bit TIFF, and created 2 light probes, one including the five at one full f-stop and the other only including the three at two full f-stops apart. There is a small difference which can be measured but is hardly noticeable. The HDRI with the one f-stop renders is less bright but has slightly better light resolution.



*Left, 5 renders at 1 f-stop; centre, 3 renders at 2 f-stops; right, difference.*

The test renders can be considered as a sun dial and show the direction of the prominent lights by the direction of the shadows cast, and their respective brightness by the shadow intensity. The left one using all five renders is 1/3 f-stop less bright than the centre one using only three renders. This was compensated by the HDRI Effect setting.

There is also a small difference of around 1° in the shadow direction; this is shown in the right picture. The green fringe comes from the left render and the red fringe from the centre render. As observed before, the difference is so small that it does not justify a 40% longer render.



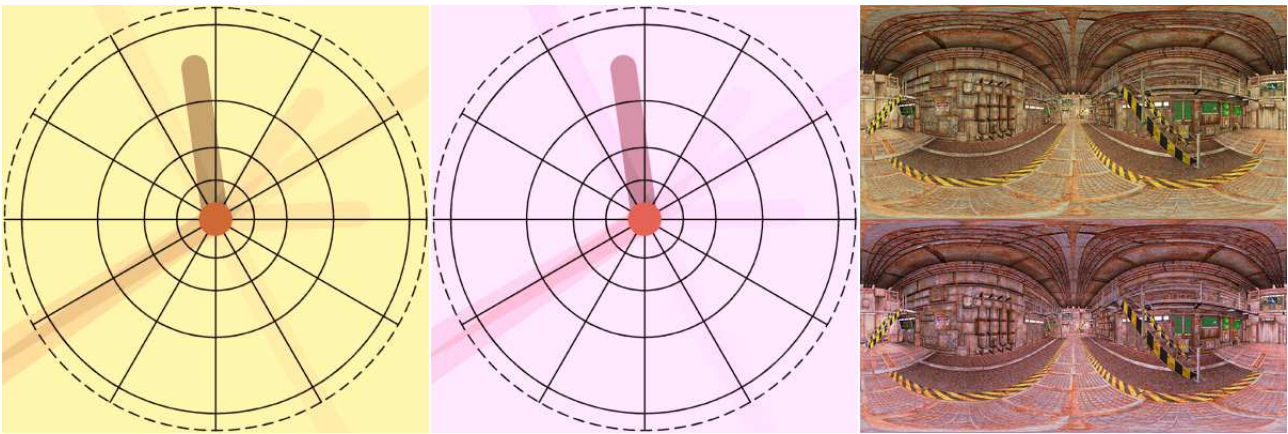
*Top row, 48 bit TIFFs used for HDRI, bottom row, the ones skipped.*

Up to now we exported the renders as 48 bit TIFF and used them to merge to an HDRI. We also proved that we can use renders with the lights in two full stop increments and we can save us the trouble to render in one f-stop increments.

What happens if we save the renders as conventional 24 bit TIFF or BMP and use those renders to assemble an HDRI? Again, I rendered in one f-stop increments and assembled two HDRIs. One takes all five renders, the other only every second.

The HDRI created from five BMPs looks nearly the same as the one made from five 48 bit TIFFs, but is less bright by about by a fourth f-stop and about one full f-stop less bright than the one made from three 48 bit TIFFs. It is acceptable nevertheless, but five renders are needed.

The HDRI made from three BMPs is twelve times brighter than the one made from three 48 TIFFs and the colours are completely wrong.



Left, five BMPs at 1 f-stop; centre three BMPs at 2 f-stops, right, the tone-mapped panoramas.

This test with the BMPs demonstrates clearly that exporting the renders as 48 bit TIFFs and using renders in two full f-stop intervals is the way to go. It gives the best result and the renders need the least of time.

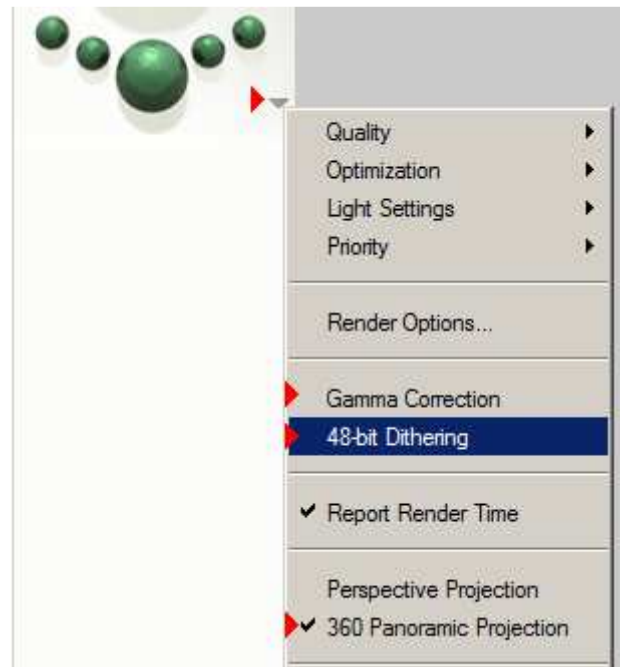
Now since this is settled for good, we will look how to export the rendered image as 48 bit TIFF. But first, we have to make sure that the render options are correctly set:

- Gamma Correction **Off**.
- 48-bit Dithering **Off**.
- 360 Panoramic Projection **On**.

Gamma correction renders the render useless for using it for assembling an HDRI because it changes the linearity of the light intensities.

48-bit dithering is cool and should be enabled if the render is saved as a conventional 24 bit image like a BMP. However, if the render is exported as 48 bit TIFF, dithering does add noise and moiré and it should therefore be disabled.

360° panoramic projection is mandatory if the Spherical Mapper is used.

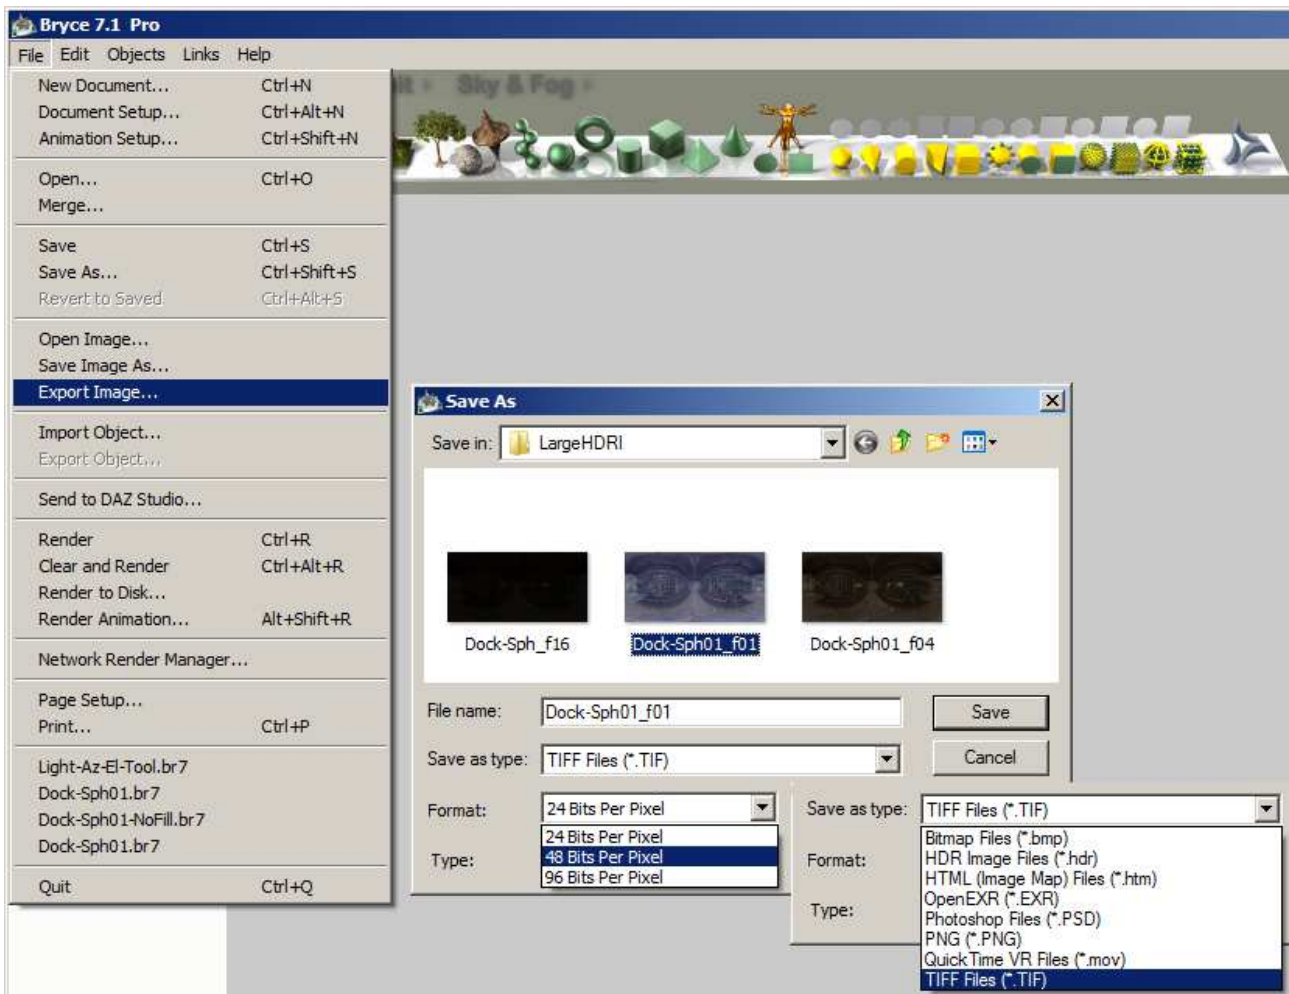


We export the render the moment the rendering process finishes as 48 bit TIFF. Why 48 bit TIFF and not 96 bit HDR or EXR? Simple – we want to merge the renders to an HDRI, an HDRI file cannot be merged to an HDRI. Only LDRIs can be merged to an HDRI. Modern DSLR cameras also can save the shots in more than 24 bit colour resolution if a RAW format is used (no matter how it is called; CRW or CR2 for Canon, NEF for Nikon, ORF for Olympus, ARW for Sony). Therefore, HDRI capable programs can read and merge 24 and 48 bit images to an HDRI.

To export the render – remember, right after rendering finished, before you do anything else – click on File, then Export Image ... this will open the Save As dialog. For Save as type select TIFF Files (\*.TIF) from the drop down list and Format 48 Bits Per Pixel.

After the next render, the file type TIFF Files is already selected, but the Format falls back to 24 Bits Per Pixel, so please pay attention, otherwise you have to render again!

Give the file a name that helps you identify which exposure the render represents. I usually add the f-stop number. It is important to know which image is which when you go about assembling the HDRI because there are no EXIF data embedded as are in a photograph.



November 2012/horo