2006-1795: CASE STUDY: USING MAYA AND MENTAL RAY FOR PHOTOREALISTIC INTERIOR LIGHTING

Marty Fitzgerald, East Tennessee State University

Case Study: Using Maya and Mental Ray for Photorealistic Interior Lighting

Introduction

This paper will explore the use of the Mental Ray renderer in conjunction with Alias' Maya to produce photorealistic interior lighting. It will present a work flow to go from a textured model, through basic traditional computer-based lighting, adding global illumination, and finally using Mental Ray's final gather and ambient occlusion to complete the lighting scheme.

This is primarily a user's guide, illustrated through a case study, to increase the realism of your renders, specifically in the lighting, rather than an in-depth guide to Mental Ray. We'll be covering the techniques in conjunction with a case study that shows the application of the concepts. This paper will provide a workflow, some real numbers along with real renders to make the numbers visibly understandable, and some indication of the time costs of the various options. It should enable anyone to be able to start incorporating these lighting techniques into their workflow.

The project for this case study was modeled and textured in Maya. Both the modeling and texturing are simple and straight-forward and could be accomplished in any modeling and texturing application. This paper deals mainly with the results from the Mental Ray renderer and, as such, they are applicable to any application using this renderer.

Throughout the paper there are references to render times. Render times are an integral part of producing these images. Given an infinite amount of time your image quality would undoubtedly increase. But who has that kind of time?

The Project

The case study will center on a visualization project involving a floor redesign of an educational building. To create the validity of the renders for the redesign, renders were first created of the floor as it currently exists. If these renders were convincing enough to be confused with current photos, then the visualizations of the redesign would be more credible. The models were completely textured to add detail and digital photography was added for views out of the windows for added realism. Throughout this paper we will be looking at two different views of the model to show examples of the specific lighting techniques.

Modeling

This project was modeled in Maya using a combination of polygons and nurbs (Figure 1 & 2). It wasn't modeled with an eye towards geometric detail, but rather simplicity in texturing and rendering. The idea was that we would attempt photographic quality through texturing and lighting rather than geometry. The model isn't heavy to work with and most objects are on display layers so they can be turned off for ease of use.

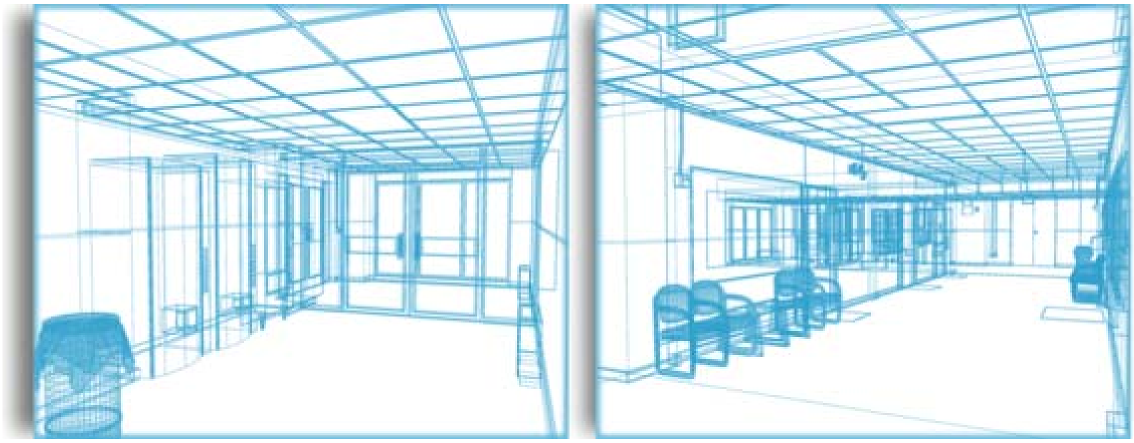

Figure 1
Model, Wireframe

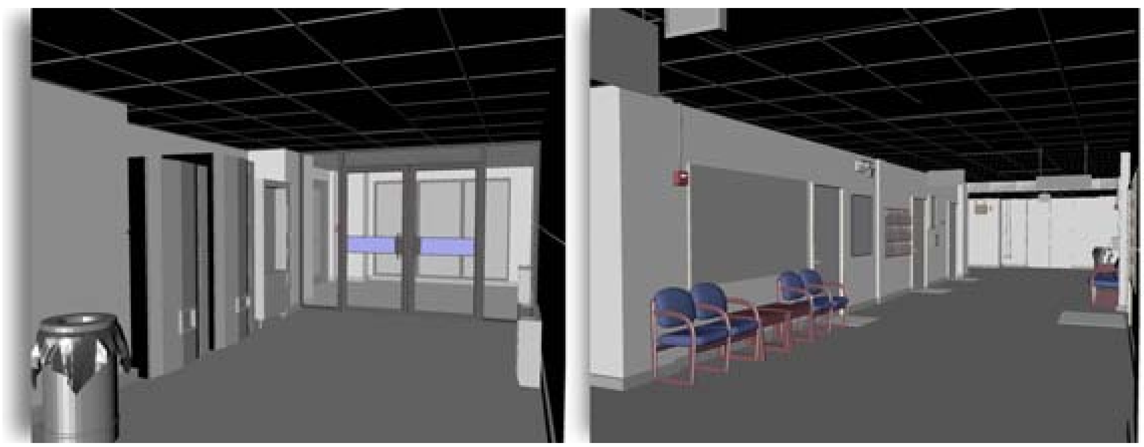

Figure 2
Model, Flat Shaded

Texturing

This model is primarily textured with planar projection of textures (Figure 3). I took extensive photos of the site in question and, through some 2D Photoshop work was able to extract color textures (and bump maps in a few cases) for the geometry. The photos had any perspective removed so that they were square to the camera, and had any baked-in lighting removed. The lighting for the textures should be supplied by the lights in the 3D scene, rather than from the conditions in which you took the picture (i.e., the texture pattern for the floor had any specular highlights from windows or lights removed so that the final renders would reflect the positions of the virtual camera and light sources). Digital photos were taken for the views out of the windows, and for some of the views into adjacent rooms.

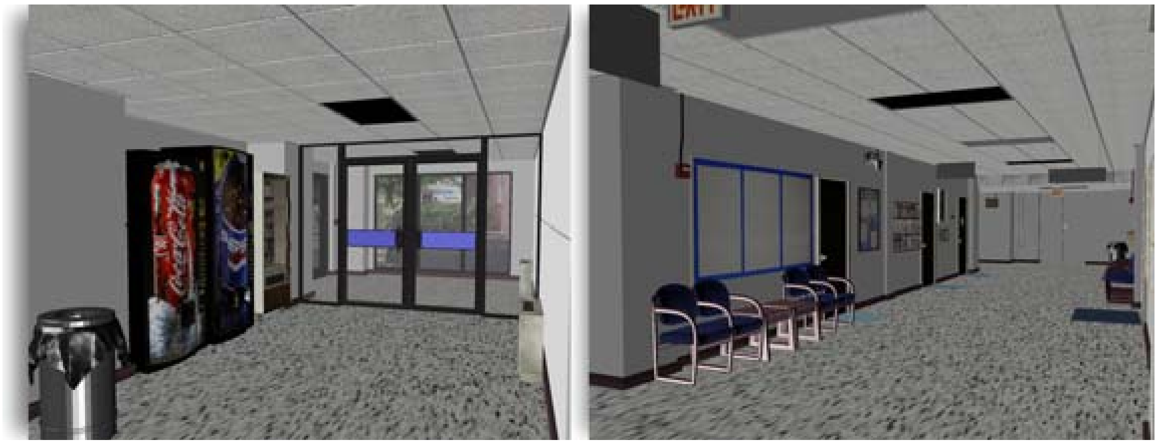

Figure 3
Model, Textured

Basic Computer Graphic Lighting

The first step in the process towards photorealistic lighting is the placement of traditional computer graphics lights: point, spot, and area. These are positioned and oriented at the locations of existing light sources (Figure 4). In this case, ceiling lights and windows.

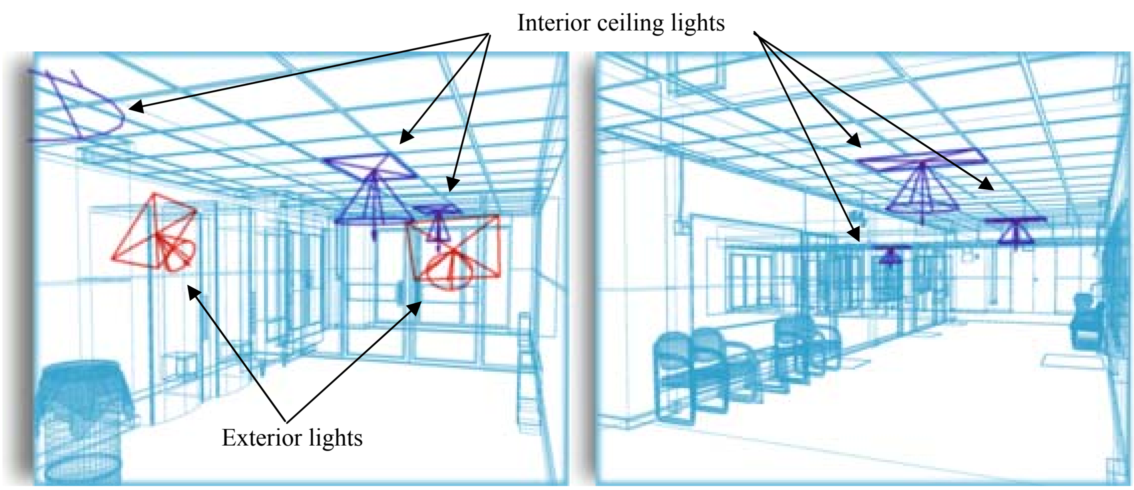

Figure 4
Light Positions

At this point, all the light in the scene is direct light. Light starts at a light source, travels to a surface and then stops. When rendered (Figure 5), the ceilings are black because no light specifically points at the ceiling and once the light hits a surface such as the floor, it stops. Some of the objects are generating their own light (the vending machines, lights, exit signs, and the outside) and have images mapped to their incandescence. They aren't adding any light to the scene, rather they require no lighting to be visible. There is no bounce or reflected light in the scene. This is the hallmark of basic computer graphic lighting.

Figure 5
Render with Basic Lighting

The next step in the typical lighting scenario is to try and simulate the bouncing light by adding more direct light sources oriented to mimic the paths of bounced light. For instance, you might put a light on the floor, pointing up at the ceiling, to emulate bounced light. While this can be effective, and certainly much better than just leaving the scene as it is, you really have to work to get the results you can achieve with the advanced computer lighting techniques.

The advanced lighting techniques covered in this paper are: global illumination, final gather, and ambient occlusion. Added on top of the lighting work that you've already done, they bring a new level of realism to your renders.

Adding Global Illumination to the Render

What is global illumination? Mental Images (makers of the Mental Ray Renderer) define global illumination as "…physically correct light simulation, which combines forward and backward ray tracing to simulate all possible light paths and light interaction from light sources to the camera..."[1] In layman's terms, global illumination is going to bounce the light that you added with the basic computer graphic lighting. So instead of light beams hitting the floor and stopping, they will bounce up to light the ceiling and continue bouncing to light some of the hard-to-reach parts of your model.

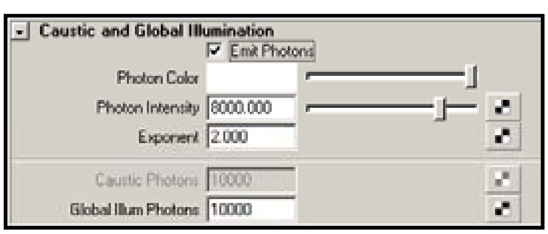

Figure 6
Global Illumination Settings Per Light

There are two places to change settings to start using global illumination. The process the Mental Ray renderer goes through for global illumination starts with each light in the scene emitting photons. This has to be user-enabled on a light by light basis (Figure 6). Simply turning the photon emission on and using the default values is a good place to start. The photons are traced throughout the scene as they bounce from object to object giving off light as they go. At some point the photons lose their energy and won't light anymore. Photon Intensity and Exponent are the primary tools for controlling how global illumination will light your scene.

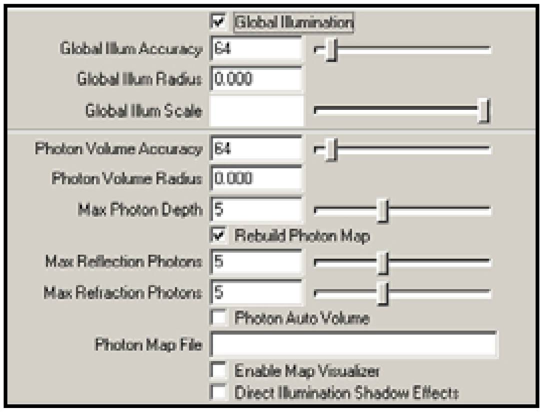

Figure 7
Global Illumination Render Settings

You must also enable global illumination in the render settings (Figure 7). The primary concern is with the first setting, Global Illum Accuracy. It defaults to 64, which is most likely going to be much too low to achieve satisfactory results, but is an okay setting to start with as it renders fairly quickly.

There are two aspects to using global illumination in your scene: (1) getting correct luminance levels, and (2) getting acceptable quality levels.

Correct luminance levels. These are a function of two settings: the photon intensity and the photon exponent. The photon intensity is the level of brightness that each photon carries as it bounces from surface to surface. Changing the photon intensity won't change the intensity of the direct (or basic) lighting, but it will change the intensity of the bounced light. In this scene there are two types of lights in play, the inside ceiling lights, and the outside light coming through the large windows. Intuitively, these light sources would have very different photon intensities. In this case study, the outside light photon intensities were set to ten times the interior light photon intensities. Figures 8-10 show different levels of photon intensity and the resulting renders. The intensity numbers are not absolute, but are relative to the geometry and lighting in this particular scene. While each scene will be different, just a few renders should be sufficient to get you in the right range for the photon intensity values.

Figure 8
Global Illumination with Photon Intensity of 1500

Figure 9
Global Illumination with Photon Intensity of 2500

Figure 10
Global Illumination with Photon Intensity of 3500

At this point, the renders will be of very poor quality, but the first step is to get the light levels correct. Since there are going to be numerous test renders at this point, there is no reason to have anything but the lowest settings for all of your renders.

How far the light bounces before it dies out is a function of the photon exponent. It defaults to 2, which is purportedly physically accurate, according to the manual. This is a very sensitive number and unlike the photon intensity values that are in the tens and hundreds of thousands, small decimal changes have significant results. Lowering this number will increase the distance a photon will travel, and therefore illuminate, before it dies out. This will effectively brighten you scene. And increasing the exponent will cause the photons to die out quicker and will darken your scene (Figures 11-13).

Figure 11
Global Illumination with Exponent of 1.5

Figure 12
Global Illumination with Exponent of 1.75

Figure 13
Global Illumination with Exponent of 2

Your final luminance levels are a balance between the direct lighting values, the intensity of the photons, and the exponent or traveling distance of the photons. You need to balance between these settings to achieve visibly correct luminance levels for your scene.

You can see the individual photons as circular blotches in all of these renders. This is what a low quality global illumination render will look like. At this point you'll determine what level most accurately creates the correct luminance level and then work to refine the look of the global illumination and smooth out the blotches.

Acceptable quality levels. Once you feel the lighting levels are correct in your scene you can adjust the quality settings until the renders look acceptable or the render-times become unacceptable. There are two places to adjust the quality of the render: the number of photons for each light and the global illumination quality level in the render settings. Both of these have an effect on the quality and on the render times and you will need to work with both. Increasing the quality levels from the default 100 into the 1000-2000 range will generally dramatically improve the quality of the render without a correspondingly dramatic jump in render times; it will increase, but not in the same ratio as the number of the setting (Figures 14-15). By adding additional photons to this you can create an acceptable render. Adding photons creates a larger jump in render times so you would use this only to the extent that you needed to, and no more.

Overall, global illumination isn't that expensive of a proposition time-wise considering the benefits. While it does slow down the render, the increase in time isn't that great compared to the increase in quality. This does take some time to set up correctly and make any number of test renders to check your luminance levels. But you would need to go through a similar process to fake the bounce light with conventional lighting and probably a greater number of renders.

Figure 14
Global Illumination Quality settings of 500

Figure 15
Global Illumination Quality Settings of 1000

Global illumination will not add correct shadowing or shading in the darker areas, however; it will add and bounce light, but not subtract any past what it will do for direct shadowing. To add that we'll use final gather and/or ambient occlusion.

Adding Final Gather to the Render

Final gather is a process for simulating global illumination that allows objects to be light sources. According to the manual, "…when used in combination with global illumination, produces the most realistic, physically accurate lighting…" Of course, when anyone in computer graphics reads "most realistic" and "physically accurate" they know that this is marketing-speak for very slow renders. This is a very render-intensive function, and while the combination of Mental Ray's final gather and global illumination does produce excellent results, the render time considerations make it an expensive option.

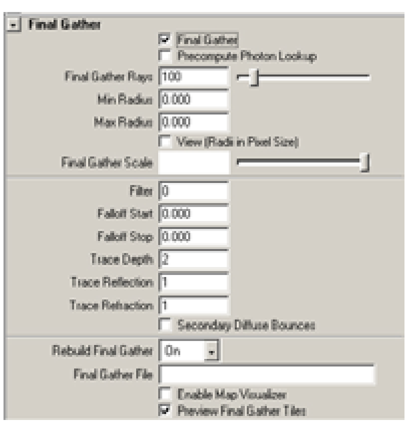

Figure 16
Final Gather Render Settings

Starting final gather is simply a matter of checking the box for it in the render settings. There are no other settings on the lights as there are with global illumination. The final gather render takes place in relation to the size of your scene so that the Max Radius and the Min Radius settings should be set to 10% and 1%, respectively, of your scene size. The number of Final Gather Rays is the main quality setting for final gather and the render times increase as this setting is increased.

All of the preceding renders have taken approximately 20 minutes for the left image with the vending machines, and approximately 40 minutes for the right image with the chairs, both at 1000 line resolution at high quality settings. Adding Final Gather brought the renders into the 2.5-3 hour range.

Figure 17
Final Gather Rays 100

Figure 18
Final Gather Rays 300

You can see from the images that adding final gather made a dramatic improvement in the believability of the render by adding darker shading under the chairs, over and under the vending machines in addition to the shadowing that was already there from the direct light. It also increased the light on the wall behind the vending machines and generally evened out the lighting extremes from the global illumination (particularly on the ceiling in the right image). When the Final Gather Rays setting is too low you will get splotches on the render. In Figure 17 you can see the splotches in the right image where the wall meets the ceiling and by the far doors, and you can see the improvement in Figure 18.

Mental Ray does have the ability to reuse its final gather calculations from render to render. So once you've gone through the time to calculate final gather from a particular camera location, you can save that file and get the benefit of final gather without the wait. This works as long as no geometry has moved and the camera is in the same position. If the Rebuild Final Gather setting is set to freeze (Figure 19), final gather will look for this user-supplied file name for pre-

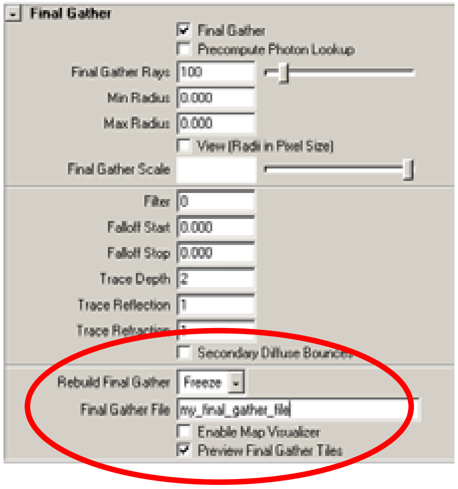

Figure 19
Final Gather Render Settings

calculated values. If the file doesn't exist, it will create it as it renders. This somewhat mitigates the render time for final gather since renders very rarely work out the first time around. Your successive renders can be done without the final gather overhead and bring you back into the 20-40 minutes range for a render. This also demands a bit of planning, however, in that the geometry, camera position and settings cannot change between renders. But shaders, colors, textures, etc. can all change.

On a side note, while these settings are producing acceptable quality still images, they are not the correct settings for animation. Final gather is calculated every frame for animation (since the camera would be moving) and would flicker or crawl at these settings.

Adding Ambient Occlusion to the Render

What is Ambient Occlusion? Ambient occlusion is a geometry-based luminance scheme that derives a value for each rendered pixel from the proximity of other geometry that occludes a 180 degree view from that pixel. Based on this occlusion, the pixel is assigned a value from 0 to 1 with 1 being no occlusion at all. The result of an ambient occlusion pass is a greyscale image that is typically used to attenuate the lighting produced by other methods (basic or advanced computer lighting). Ambient occlusion renders are typically white with darker details.

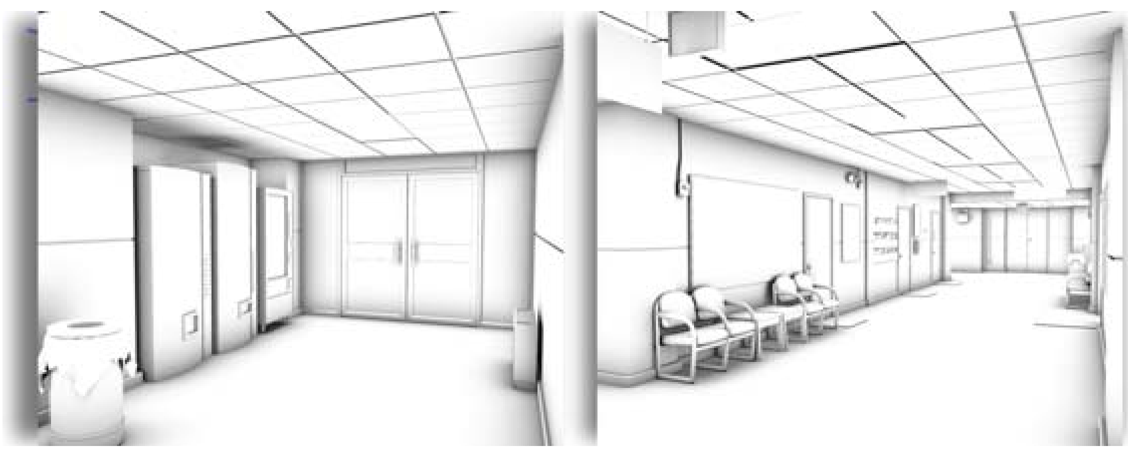

Figure 20
Ambient Occlusion

Settings for ambient occlusion are related to the scale of your model. The primary control is the Max Distance setting which determines how close a piece of geometry has to be to rendered pixel to occlude it. There is also a Samples setting to determine the quality. If the render is looking too splotchy or grainy, you can up the samples setting to smooth it out (and increase the render time).

Looking at Figures 21-23 you can see that as the Max Distance setting increases, the likelihood that a rendered pixel will be occluded increases. The darker areas extend farther and farther, say, from the base of the walls as the Max Distance increases. (There is some flexibility in rendering with a higher number for the Max Distance in that you can adjust the values down in Photoshop after the render while you can't accurately adjust them up.)

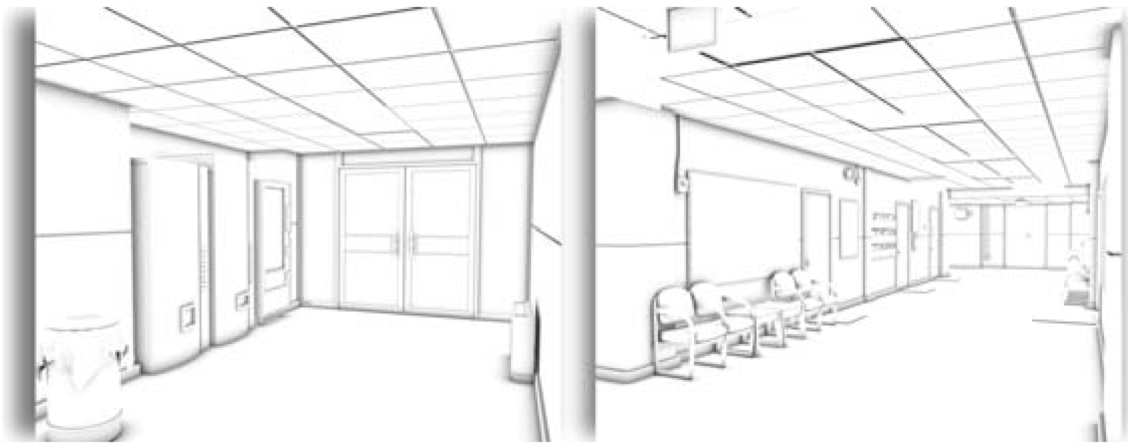

Figure 21
Max Distance = 1

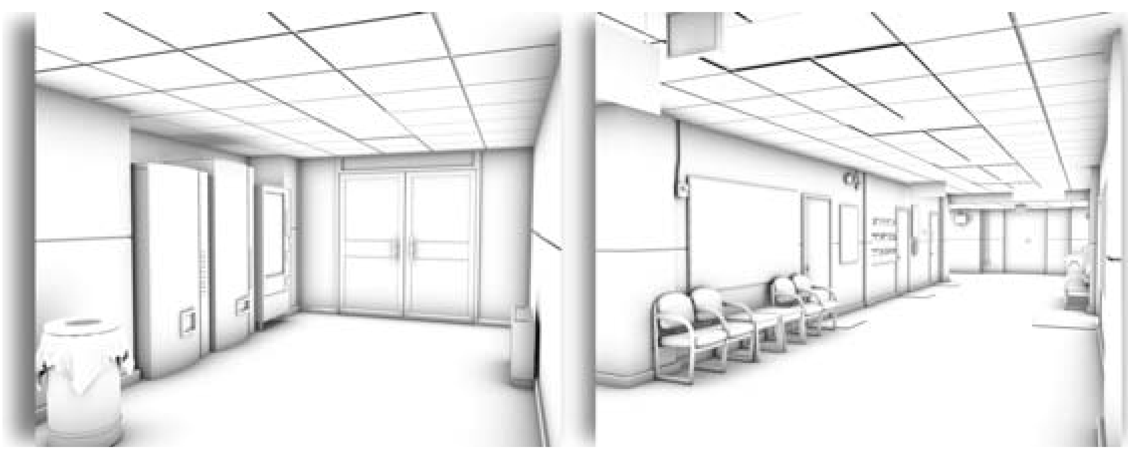

Figure 22
Max Distance = 3

Figure 23
Max Distance = 6

As you can see, ambient occlusion renders are primarily white, with grey to black edges detailing objects in your scene. It is typically layered on top of your render (in Photoshop or a similar program) in a layer with a layer mode set to “Multiply.” Everything that is white (or 1) in the ambient occlusion pass remains unchanged in the final image, while everything that is black in the ambient occlusion pass (or 0) is black in the final image. Grey scale values in the ambient occlusion pass affect the final image accordingly. Figures 24-25 show what an ambient occlusion pass can add to a render. You can notice the shading: under the chairs, where the walls meet the ceiling, under the vending machines, where the walls meet the floors, etc. Where ever there is dark shading in the ambient occlusion pass, you see darker shading in the final render.

One advantage of using the ambient occlusion pass is that once your camera position has been identified, you can render the pass at any time. You don’t need to wait until everything is textured, lit, etc. Final gather is a part of your final render and is rendered along with everything else, whereas the ambient occlusion is a separate render pass and is composited in a separate process. This makes it advantageous for fitting it into a production pipeline.

Ambient occlusion and final gather are not the same thing. But they seem to work somewhat similarly when combined with global illumination in that they both appear to subtract light (or add shadowing) in parts of your scene. Ambient occlusion doesn’t add any luminance anywhere, or smooth out some of the lighting from the global illumination as final gather does. They can both be used together, or separately. The ambient occlusion pass has the big advantage in that tweaking its settings is very quick, and rendering is quick as well, so there isn’t a big time overhead in rendering this pass. This pass will add also help significantly with a traditional render done with direct lighting only. When this pass is layered over your other renders in Photoshop you adjust the amount of the effect to suit by adjusting the transparency of that layer.

Figure 24
Render with Global Illumination only

Figure 25
Render with Global Illumination and Ambient Occlusion

Quick Review of the Process

Since the actual steps towards producing this type of work can get lost in the descriptions and pros and cons, below is a quick, step-by-step guide to the process:

1. This process starts with a 3D model that is textured.
2. Objects that generate their own light are textured accordingly in the incandescence channel of their shader.
3. Basic lights are added in positions and orientations that correspond with real light sources in the scene and set to cast shadows.
4. Global illumination is added.
    a. Photons are turned on at each light source.
    b. Global illumination is turned on in the render settings.
5. The photon intensity and exponent settings on the lights are adjusted to give the correct illumination of the scene.

6. The quality of the global illumination is adjusted to bring the render quality up to an acceptable level and get rid of the blotches. This is done by increasing the number of photons in each light, and increasing the global illumination quality setting under the render settings.
7. Possibly add final gather to the render.
    a. Turn final gather on in the render settings and increase the number of rays to improve the quality.
        i. You can set the Rebuild Final Gather setting to Freeze if you want to pre-generate the final gather calculations.
8. Possibly add ambient occlusion to the render.
    a. Once the camera and geometry are set you can render the ambient occlusion pass at any time.
        i. The main setting is the distance setting which refers to how far a piece of geometry can be and still occlude a rendered pixel.
        ii. There is also a quality setting that will smooth out the graininess.
        iii. Once this is rendered it is typically layered over the global illumination render in Photoshop on a layer with its layer mode set to Multiply.

Conclusions

People are constantly exposed to some of the best (and worst) computer rendered images in television, movies, and print. Unlike the engineering and science that goes into those images, the average person is, to some degree, an authority on whether it is photo-real or not. They might not be able to put into words exactly what is right or wrong with an image, but they feel comfortable and even enjoy judging its quality.

While it is certainly possible to create these photo-realistic renders using traditional computer graphic lighting methods, you can increase the quality while reducing the time necessary to generate these renders by using some now commonly-available tools — global illumination, final gather, and ambient occlusion — in a commonly-available renderer: Mental Ray. These renders may not be 100% physically accurate, but they extend the believability of an image and take it from “computer-generated” to “photographic.”

There is most certainly a learning curve in using these techniques, but it’s not that extensive compared to the improvement in your renders. Learning the techniques is much facilitated by using very simple geometry, textures, and lights to eliminate any extraneous influences on a render so that you can focus on lighting and render settings. This will quickly give you a feel for how global illumination, final gather, and ambient occlusion work. These techniques are probably not best debuted in a project with very tight deadlines, but rather in a setting where there is some time for revision and refinement.

As this is a case study, I should note where the project currently stands: the project is past its due date, and has yet to start. In spite of all the render time spent on this project, the slowest element in the pipeline is the decision process.

Figure 26
Final Renders

Bibliography

1 http://www.mentalimages.com/2_1_1_technical/index.html

2. Hayden Landis 2002, "Production-Ready Global Illumination" Course 16 notes, SIGGRAPH 2002.

3. Janne Kontkanen and Samuli Laine 2005, "Ambient Occlusion Fields" Proceeding of the 2005 symposium on Interactive 3D graphics.

4. Darren Brooker 2002, "Essential CG Lighting Techniques" Focal Press.