Low-Light-Level Simulations and Extended Area Source Modeling

Emmett Ientilucci

October 11, 2001

Contents

List of Figures

List of Tables

1 Introduction

This report exemplifies recent upgrades to RIT's DIRSIG model, the latest version of which includes man-made secondary sources and exoatmospheric sources such as the moon and starlight. These new enhancements enable the user to generate a wide variety of low-light-level scenarios. The completed task at hand, however, extends DIRSIG's capabilities even further. These new improvements enable DIRSIG to directionally shape light sources, simulate extended sources, and directly view point and extended sources. Additionally, the model now incorporates lunar scattering phenomenology. This report also includes a low-light-level validation proposal that rigorously evaluates all the parameters needed to do an end to end validation of DIRSIG. Finally, a section has been included that investigates the spectral distributions of various sources, some of which are found in typical low-light-level scenes.

2 Simulation of Point Sources

This section details the additional functionality added to point sources in DIRSIG. This improvement takes the form of beam shaping and the ability to directionally point a source. Previously, the SIG model incorporated point sources, however, the distribution of such sources was only isotropic. The improvement comes in a way that the user can, not only specify the location and distribution of a source, but can now specify both the direction and shape as well.

2.1 Isotropic Point Sources

Figure 1 shows a low-light-level image captured with an image intensified CCD. The light source in the image exhibits a gradual falloff in illumination as one scans across the beam profile on the ground. Previously, to simulate the diffraction and internal reflections from such a light box, a series of 5 point sources (1 inch apart) had to be used (see Figure 2). The result was a "shaped" light field with noticeable artifacts. Clearly the aperture edges can be seen on the ground as generated by each point source in the housing.

2.2 Non-Isotropic and Directional Point Sources

The solution for the above scenario is to factor in a weighting function that modulates the light distribution as a function of view angle. This is

Figure 1: Low-light-level image showing a light source in a small housing.

Figure 2: a) DIRSIG LLL image showing 5 stacked point sources and b) side view of source in housing.

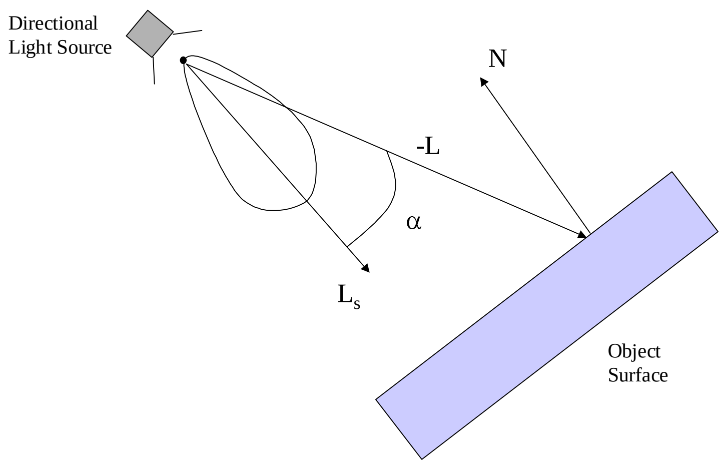

Figure 3: Directionally shaped light source.

accomplished with a simple trigonometric function raised to a power. Here we assume that for a directional source, the light intensity in a particular direction, given by the angle α, is:

$$I_s \cos^m \alpha \tag{1}$$

where m can be thought of as an exponential shaping factor. The angle α is between L, the direction of the point on the surface that we are considering, and L_s, the orientation of the light source (see Figure 3). The value of I_i that we use in the shading equation is then given by:

$$I_i = I_s(-L \cdot L_s)^m \tag{2}$$

$$-L \cdot L_s = |-L||L_s| \cos \alpha \tag{3}$$

A plot of this function with various values for m can be seen in Figure 4. The plot increments α from 0 to π, where a values from $\pi/2$ to π get wrapped around, tracing back to zero. From this plot, one can clearly see that as m gets larger the cosine term (or modulation term) becomes more confined to the zero axes thus shaping the intensity distribution.

It is also seen that this technique could resolve some diffraction issues involving aperture sources. That is, the user could compensate for aperture effects by shaping the light field. Additionally, this shaping factor is more

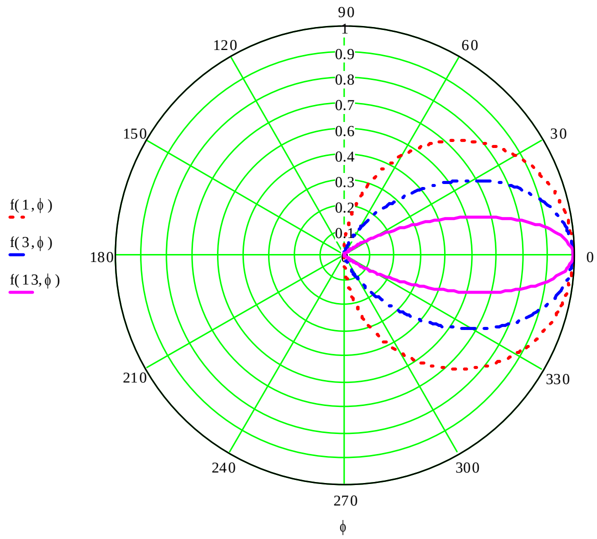

Figure 4: Polar plot of intensity shaping function

appropriate for modeling focused sources such as streetlights, headlights, and desk lights. If one wishes to still model an isotropic source, the value of m is simply set to zero.

The effects of the shape function can clearly be seen when implemented on DIRSIG imagery. Figure 5a shows an image simulating a shaped light source by using 5 point sources, each placed 1 inch apart. Figure 5b shows the same image with the new shaped source. The results are more natural looking and closely resemble those found in the truth image of Figure 1.

3 Simulation of Extended Sources

One of the main disadvantages of ray tracing is its ability to simulate wide area or extended sources. The previous sections dealt with tracing rays to a 3-dimensional point in space, which contained a point source with a known spectral distribution. This tracing of rays is easy to accomplish in a ray tracing environment since one ray can trace directly to one source. The problem is compounded, however, when the point source is no longer a single 3D point in space but an infinite number of points that make up an extended area source. Tracing to this many points is certainly computationally prohibitive. There are other methods that circumvent this problem,

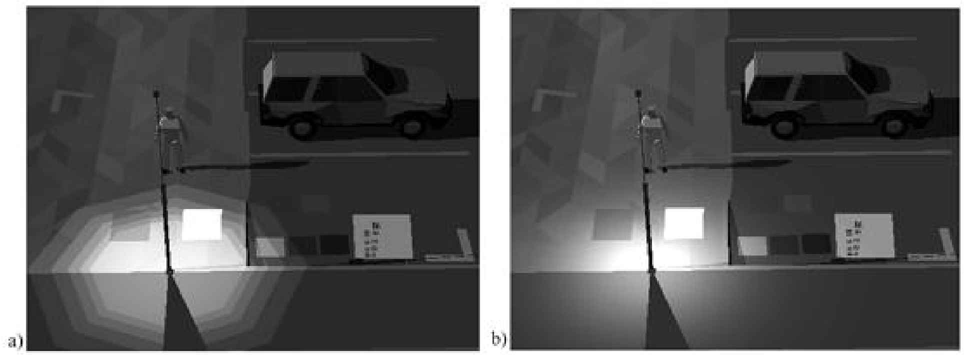

Figure 5: a) Light with 5 point sources and b) light with 1 shaped point source.

such a radiosity, but they are very computationally expensive and, more importantly, impractical for our application in remote sensing. A solution has been developed, however, that samples an area source in such a manor as to generate an approximation to what the radiance reaching a target would be.

3.1 Global Illumination

In computer graphics, global illumination is the term given to models which render a view of a scene by evaluation the light reflected from a point x taking into account all illumination that arrives at a point. That is, we consider not only the light arriving at the point directly from light sources but also all indirect illumination that may have originated from a light source via other objects. Partial global illumination models do exist, such as ray tracing and radiosity; however, they both simulate only a subset of global interactions. Ray tracing tends to account for (perfect) specular interactions while radiosity accounts for (perfect) diffuse interactions.

3.1.1 Ray Tracing, Radiosity, and Alternative Solutions

Ray tracing traces light rays in the *reverse* direction of propagation. That is, from the eye back into the scene towards the light source. To generate a two-dimensional image plane projection of a scene using ray tracing, we are only interested in the light rays that end at the sensor or eye point. Therefore it makes sense to start at the eye and trace rays out into the scene. Ray tracing is thus a *view-dependent* algorithm. Radiosity implements diffuse-

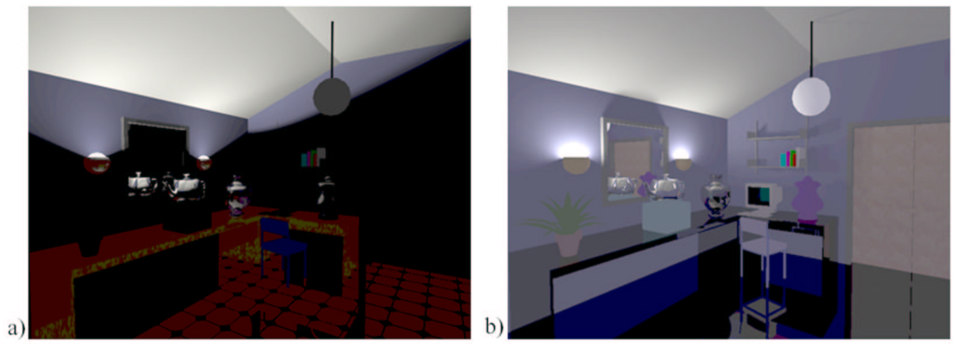

Figure 6: a) Ray traced and b) radiosity rendered images. Image b) has also been ray traced for specular objects.

diffuse interaction. Instead of following individual rays, it is the interaction between patches (or polygons) in the scene that is considered. Thus the solution is *view-independent* and consists of a constant radiosity for every patch in the scene. To see examples of these algorithms consider Figure 6. Most of the global interaction in the scene is diffuse-diffuse (*i.e.,* indirect lighting from the upward facing lights on the wall). The ray tracing approach lacks the ability to quantify multiple bounces (Figure 6a). The radiosity approach solves the multi-bounce energy equilibrium and effectively accounts for infinite bounces (6b).

To account for all possible sources of photons with ray tracing, an infinite number of rays would be needed. Furthermore, artifacts can be incurred if a limited numbers of rays are used. The radiosity solution approaches this differently by computing reflected photons on a per-facet basis. To fine tune the solution, facets containing large flux gradients need to be sub-sampled. This means that current DIRSIG scenes could grow by several orders of magnitude in size (facet count). Additionally, this method would require a significant amount of code development to make it practical. In general, radiosity solutions are used for indoor scenes with thousands or tens of thousands of facets. Typical DIRSIG scenes contain hundreds of thousands of facets.

There is yet another solution that involves the use of light maps. These are created for a single facet surface to describe the distribution of incident light. They could be incorporated into DIRSIG in a similar manner as texture maps (*i.e.,* indexing into a map) and are pre-calculated using any rendering method (*e.g.,* radiosity). However, the creation and storage of maps is time and storage expensive again making them impractical for this application.

To overcome this limitation we have developed and implemented a solution that takes advantage of the specific attributes of the DIRSIG ray tracing approach.

3.2 Methodology for Ray Tracing Extended Area Sources

A facet in the scene is flagged as a source by the material type assigned to it. The facet vertices are used to compute a surface or direction normal for the source using a "right-handed" or clockwise rule. The DIRSIG ray-tracing based treatment of extended sources uses an area sampling technique that relies on computing the exposure or "shape factor" of a given source to a point in the scene. This is accomplished by tracing a series of rays to the source to estimate the apparent occlusion of the source by other objects.

3.2.1 Pre-Render Calculations

To simplify the rendering time calculations, each source has a set of sampling points, M, that are automatically generated during the initialization of the model. These 3D points are generated randomly using a uniform distribution within the plane of the source and within the vertices. The points will be used later in the rendering process to estimate the obscuration of the source. The maximum number of sample points can be controlled through the `MAX_SAMPLE_POINTS` variable in DIRSIG.

The number of sample points, M, for any given source is computed by multiplying the source area, A, by the user controllable `SAMPLE_AREA_DENSITY` variable. The estimated area of the source is computed under the assumption that the source is either a triangle (for 3 sided sources) or a rectangle (for 4 sided sources). A summary of these steps, for each area source, is illustrated in Figure 7.

3.2.2 Pre-Render Calculation of Source Area

For rectangular as well as triangular shaped sources, the aspect ratio of the source (*i.e.,* length vs. width) is important and needs to be considered. For a rectangular source, the average of the two opposite edge pairs (l and w) is computed and the area is computed as the $l \times w$ (see Figure 8). For triangular sources, a simple average of the leg lengths will not work. That is, an equation of the form $A = \frac{1}{2}(l_{avg})^2$ will over estimate the source area for shapes with large aspect ratios (*i.e.,* large differences between base and

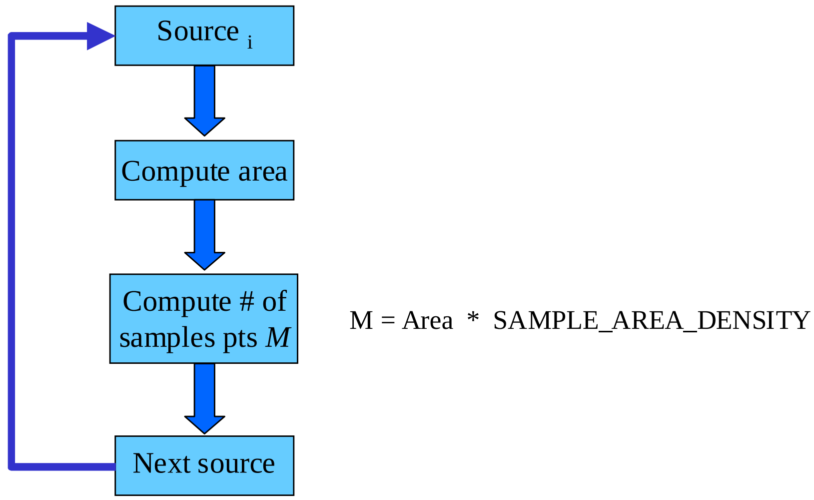

Figure 7: Flowchart of pre-render calculations performed on extended area sources.

height). A solution that over estimates the area to a lesser extent is to use the minimum and maximum length legs in the equation (see Figure 8).

Again, both approaches will over estimate the area for sources with large aspect ratios, however, doing this will only increase the number of sample points, M, generated for the source and improve the accuracy of the modeling.

3.2.3 Rendering

During the rendering process, the radiance from a given source at a given location in the scene is computed using an area sampling method. The contribution from a given source at a given target point is dictated by the relative target point-to-source geometry, the amount of the source that is not obscured at the target point, and some governing radiometry. A flowchart illustrating this algorithm can be seen in Figure 9

When rendering a pixel (*i.e.,* hit point), each source in the scene is taken into account. For a given hit point, we compute the solid angle, Ω, between the pixel and $source_i$ in the hemisphere above. This angle is computed in a similar manor to that of the source area, which was performed in the pre-rendering stage. Some of the angles used in the calculation are illustrated in Figure 10. The solid angle is then computed as follows:

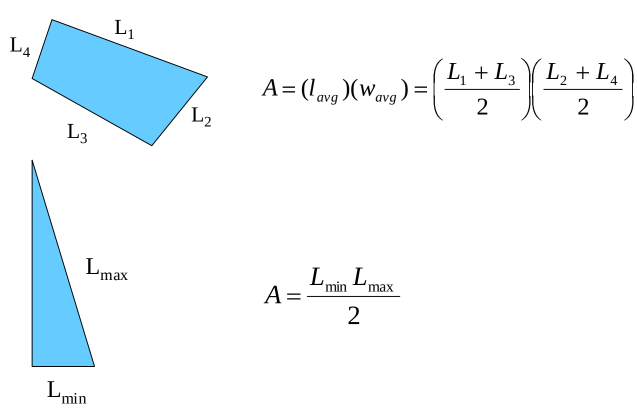

Figure 8: Area calculation for 3 and 4 sided sources.

$$\Omega_i = \left(\frac{\phi_1 + \phi_3}{2}\right)\left(\frac{\phi_2 + \phi_4}{2}\right) \qquad \text{4 sided source} \tag{4}$$

$$\Omega_i = \left(\frac{\phi_{min}\phi_{max}}{2}\right) \qquad \text{3 sided source} \tag{5}$$

Where ϕ is an angle formed by the sides of the source to the hit point. If the solid angle is smaller than the (user configurable) `SOLID_ANGLE_THRESHOLD`, then the contribution from the source is assumed to be negligible and we move on to the next source. When the solid angle is *larger* than the threshold, we calculate a number of *source sample,* N, points out of the larger set, M. N is computed by multiplying the apparent solid angle, Ω, of the source by the `SOLD_ANGLE_DENSITY`. That is

$$N = (\Omega)(\texttt{SOLID\_ANGLE\_DENSITY}) \tag{6}$$

The value of N will be different at each hit point because the solid angle varies as a function of hit point location. Furthermore, using the `MAX_SOURCE_SAMPLES` variable can put a limit on the value of N (to be described at the end of this section). Each of these *source sample* points are randomly selected and ray-traced to in order to estimate the amount of the source that is obscured (see Figure 10). If a ray hits an object on its way to the source, the transmission of that object is taken into account. The object may be opaque or transmissive (*e.g.* tree leaf). This random sampling and obscuration checking of the source is what yields the soft and gradual shadows found in scenes containing extended area sources.

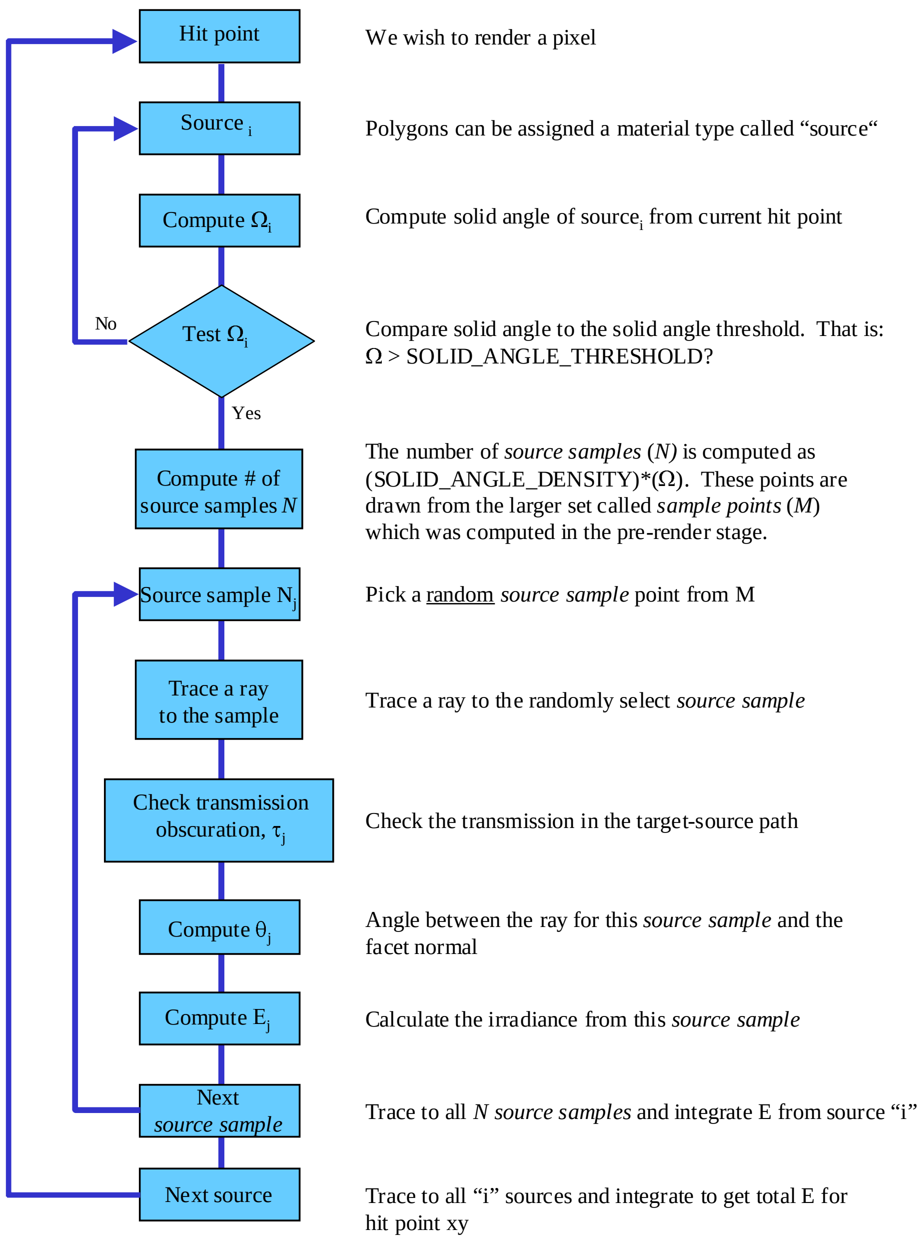

Figure 9: Flowchart illustrating how a pixel is rendered when influenced by an area source.

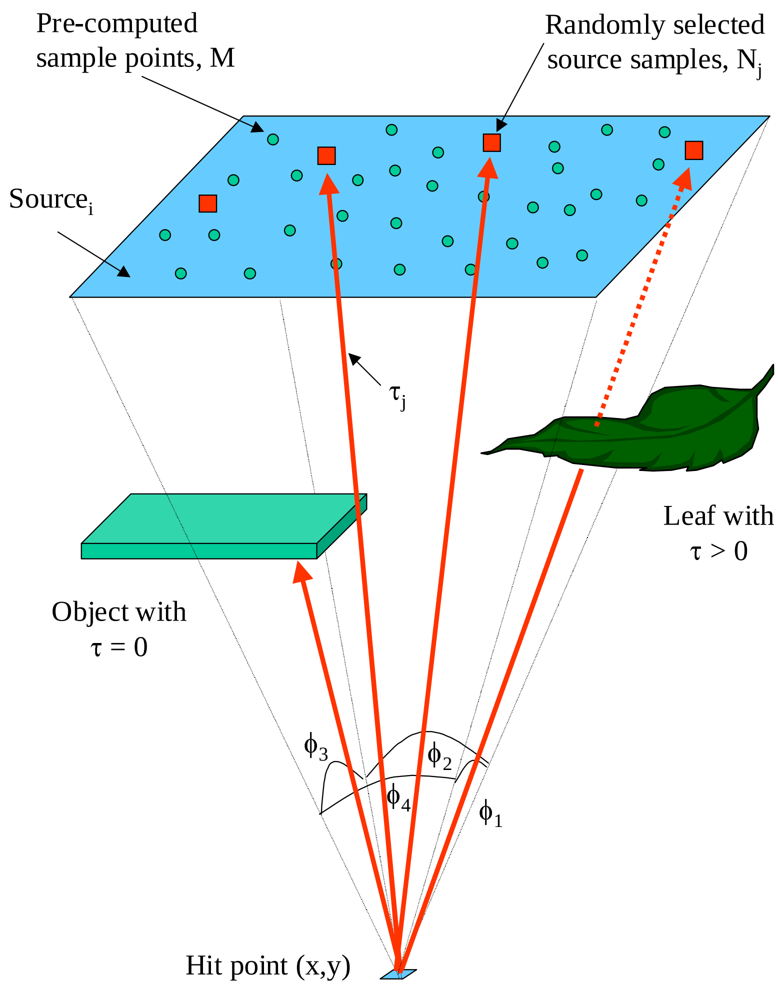

Figure 10: Ray-tracing to an area source. The area source has a pre-computed set of sample points, M. Once at a hit point, we trace N rays to the source based on the solid angle calculation (which will be different at the next hit point). In this case 4 rays are traced but only 3 actually make it to the source. One of the 3 rays is attenuated due to transmission effects thereby reducing its irradiance contribution to the current hit point.

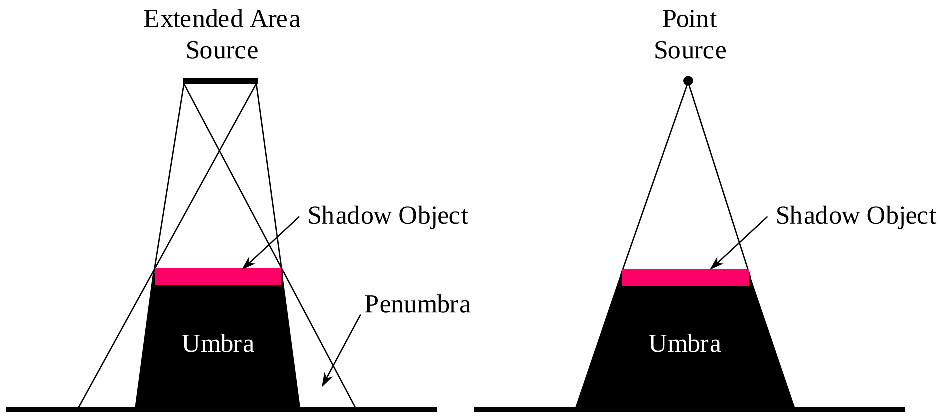

Figure 11: Illustration of shadowing due to extended area source.

Figure 11 illustrates why the random sampling and obscuration checking are so important in the ray-tracer. The shadow volume, or *penumbra*, behind the object lit by an extended area light source (in contrast to a point source) doesn't have sharp boundaries. This is caused by the fact that each point in the boundary area is only *partially shadowed*. The *umbra* is that part of the shadow that is completely cut off from the light source. This is the only shadowing that incurred form using point sources. With an extended source, however, a penumbra is generated. This is the area that receives *some* light from the source. A penumbra surrounds an umbra and there is always a gradual change in intensity from a penumbra to an umbra.

Continuing on in the rendering process, following the flowchart in Figure 9, we next compute the angle, θ_j formed by the facet normal and current *source sample*, N_j. This angle is needed in the irradiance calculation. Lastly, the irradiance at hit point xy for the first source, for example, is computed as

$$E_1(\lambda) = \frac{1}{N} \sum_{j=1}^{N} \frac{L(\lambda) \cos^2 \theta_j \, A}{R_j^2} \, \tau_j(\lambda) \cos^m \alpha \tag{7}$$

Where L is the spectral source radiance, A is the area of the source, R is the distance to *source sample* N_j, τ_j is the spectral transmission, α is the angle formed by the lighting vector, L_s, and the direction to the hit point, and m is the shaping factor.

This same calculation is performed on all subsequent sources in the scene. Additionally, DIRSIG calculates what the source contribution in the spec-

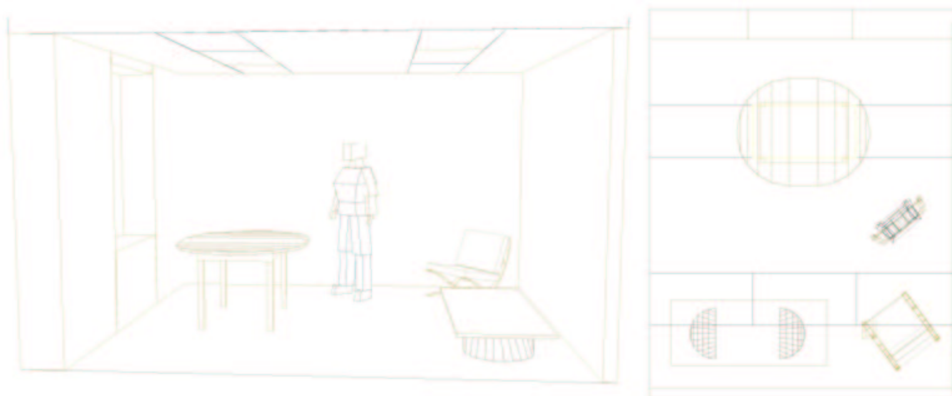

Figure 12: CAD drawing of side and top view of indoor office scene.

ular direction (with BRDF turned off, for example) would be as well. The final integrated irradiance for the current hit point at position xy, excluding the specular direction, is simply the sum of the individual contributions from the i^{th} source. That is

$$E(\lambda) = \sum_{i=1}^{k} \frac{1}{N_i} \sum_{j=1}^{N_i} \frac{L_i(\lambda) \cos^2 \theta_{ij} \, A_i}{R_{ij}^2} \, \tau_{ij}(\lambda) \cos^{m_i} \alpha_i \tag{8}$$

where k is the number of sources. To see how this relation is derived, refer to Appendix A on page 68.

3.3 Extended Area Source Results

To test the performance of ray tracing to large area sources, an indoor test scene was constructed that contained two ceiling lights (see Figure 12). The lights were typical indoor fluorescent office panel lights. This scene also contained various objects, such as chairs and tables, to check for correct shadowing. One object in particular, a large table, was placed directly below one of the panel lights. With this geometry differences between tracing to a point source and extended area source should become evident.

3.3.1 Ray-Tracing to a Directionally Shaped Point Source

As a starting point, we first illuminate the test room with a *directionally shaped source*. This can be seen in Figure 13. Here we see the effects due to an isotropic point source (*i.e.*, shape=0) and a source with a shape value

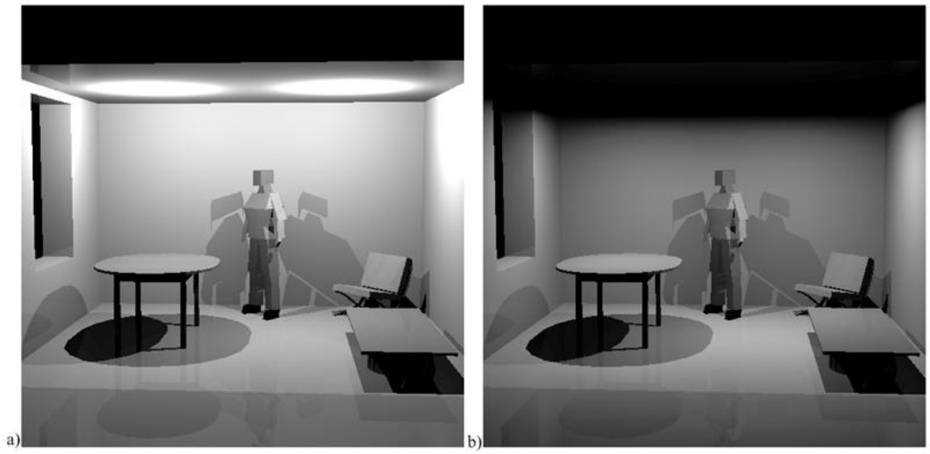

Figure 13: Indoor office scene with two lights. a) Shape=0 (*i.e.*, pt. source) and b) shape=1, both scaled the same.

of one. It is first noticed that the sharp shadows on the ground (from the table) are inherently due to the nature of a point source. The hit points on the ground either see the source or they don't. This creates the sharp transition on the ground. However, it is common knowledge that this type of shadowing does not occur in the real world, especially when a panel light is being used as the source of illumination. The shadow pattern should be soft and gradual, unlike what is seen here. This softness is created because an extended area source can be thought of as an infinite number of point sources over the given source area. Thus, a single point on the ground sees many points, or a fraction of the source. As the angle to the source decreases, the area fraction also decreases diminishing the intensity at that point.

3.3.2 Radiosity Solutions for an Extended Area Source

As a means of comparison, a series of radiosity solutions were created using a well-established algorithm embedded in a commercially available software package called Lightscape (see Figure 14). Lightscape is an advanced lighting and visualization application used to create accurate images of how a 3D model of a space, or object, would appear if physically built. It also uses a physically based interface for defining lights and materials. Lightscape is more of a photometric model only taking into account that, which pertains to the visible part of the electromagnetic (EM) spectrum. DIRSIG, how-

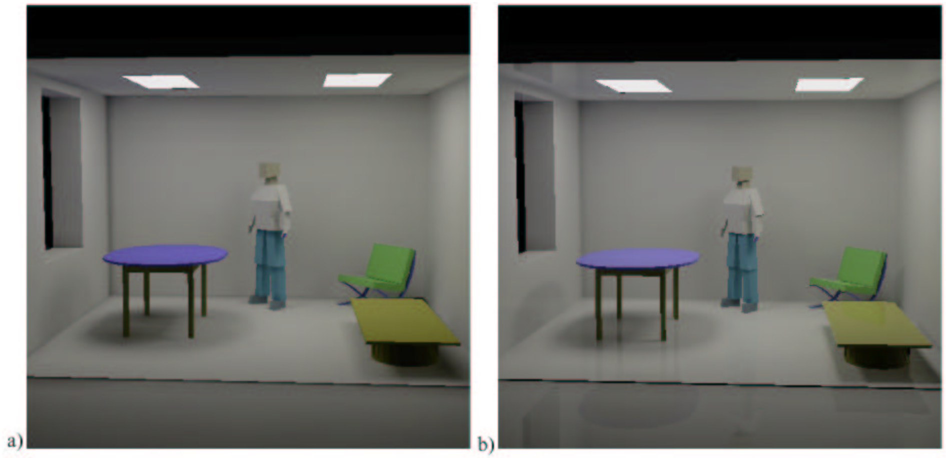

Figure 14: Radiosity solution for indoor office scene. a) Radiosity only and b) radiosity with ray-trace for specular objects.

ever, is a full-blown radiometric model that simulates any part of the EM spectrum from 0.38 to 20 μm in an arbitrary number of bands. There are other significant differences between the two models, however, these details are beyond the scope of this report.

The solutions presented exhibit much of what would be expected from a room lit with real ceiling panel lights. The shadows on the floor and wall due to the table and human are softer than that which was created using point sources. Since a radiosity solution is best suited for the test room (because of the diffuse-diffuse interactions) we use the radiosity results as a set of "truth images". These can then be used to compare and contrast the results found using the ray tracing technique described earlier.

3.3.3 Ray-Tracing to an Extended Area Source

Figure 15 shows two DIRSIG images, one created using two point sources, the other using two extended area sources. The source spectral radiance used was that obtained using an ASD spectrometer. The image was rendered from 0.4 to 0.7 μm.

The differences in the two images are fairly significant. It is first noticed that the point sources exhibit a spread of energy across the ceiling. This is expected from an isotropic point source. With the extended area sources, however, we see the energy is confined to the facet or polygon in which

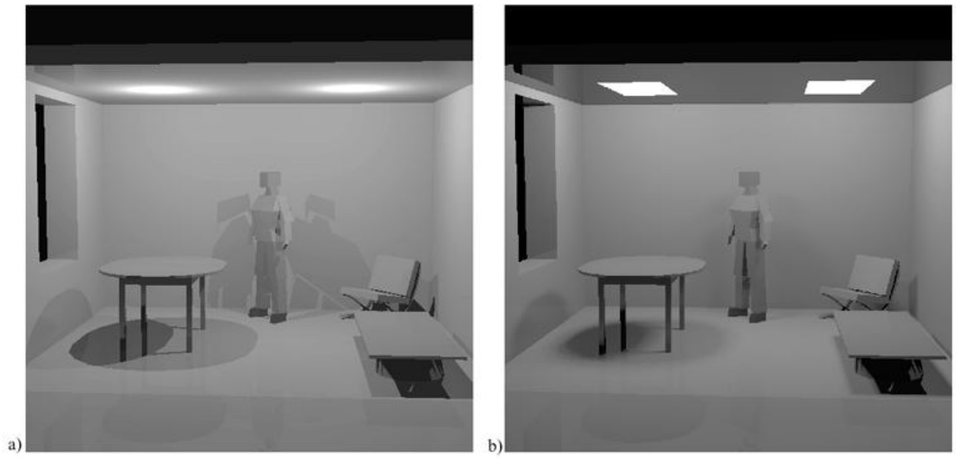

Figure 15: DIRSIG images generated using a) two point sources and b) two extended area source.

the source was defined. These two techniques also produce vastly different geometric shadowing results. The point source produces very sharp shadows (as expected) while the extended area sources produce the penumbra and umbra areas discussed earlier. This is very evident when looking at the tables and human shadow on the back wall. Another interesting phenomena to point out is the illumination magnitude on the ceiling. The ceiling using the extended area sources is much darker than the one containing point sources. The reason being is that the ceiling in the point source case is being *directly* illuminated while that of the area source is being *indirectly* illuminated. This result, in the extended area case, illustrates the recursion of the ray-tracer. That is, the ceiling, in reality, can not directly see an area source. Most of its illumination comes from the walls and floor below. If the ray-tracer could not trace to background objects after the initial hit point, then the ceiling would appear black and would never be illuminated. More details on the recursive nature of the ray-tracer to follow.

We now compare DIRSIG results to that produced by Lightscape. These two images can be seen in Figure 16. The Lightscape image was produced using a radiosity solution. The image was then ray-traced (in Lightscape with a recursion of two) for specular objects. The DIRSIG image was produced using the ray-tracer and the algorithm explained earlier. Scaling aside, the results are remarkably similar. The ceiling in the Lightscape image is brighter because the radiosity technique is more adept at the handling

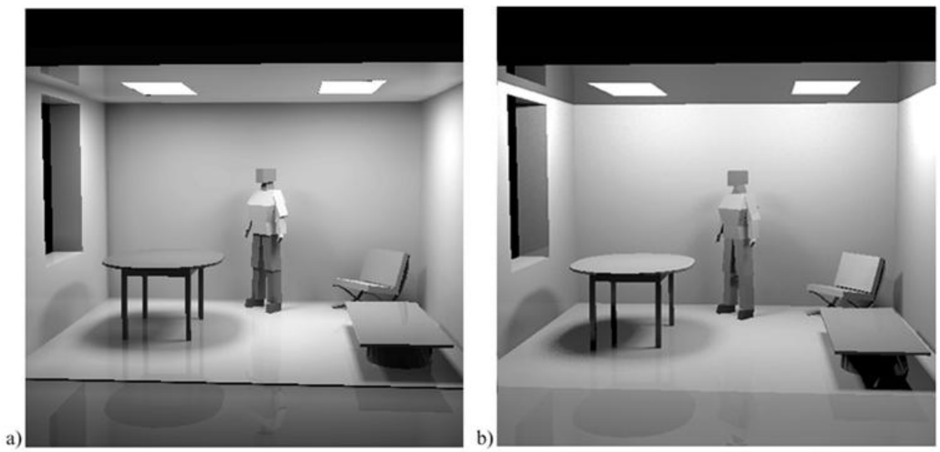

Figure 16: a) Lightscape image created using a radiosity and ray-trace solution. b) DIRSIG image created using ray-tracing only.

diffuse-diffuse interactions. As previously mentioned, the bulk of the illumination on the ceiling is caused by this diffuse-diffuse radiational exchange. DIRSIG does do an adequate job of simulating this phenomenon, even with a recursion of two.

The shadowing and specular highlights in both the Lightscape and DIRSIG solutions are identical. The facets or polygons in both images were set to have the same material attributes, including reflectivity and specularity. It is seen that the geometric shadowing in the DIRSIG image looks exactly that same as that produce by the radiosity solution. The DIRSIG area sampling technique for extended area sources has successfully produced that which is typically seen in radiosity-type solutions.

We now turn our attention to the importance of recursion in the ray-tracer. Figure 17 shows two DIRSIG images with different recursion values. The image with a recursion of one simply checks for the source contribution at the given hit point then moves on to the next hit point. This means that the ceiling will have a radiance of zero since it cannot see any of the sources. Similarly, a majority (two-thirds) of the shadow area under the coffee table is zero as well for the same reason. If we implement a recursion of two, we see that the ceiling and coffee table shadows appear much different. When a hit point lands on the ceiling, DIRSIG will trace again (hence recursion of two), in the specular direction. It will trace in more directions if the BRDF model is turned on. However, for now we concern ourselves with the

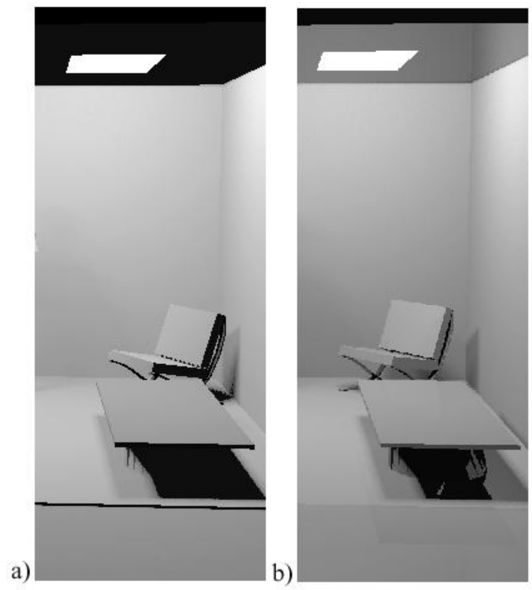

Figure 17: DIRSIG images showing the importance of recursion. a) Has a recursion of one while b) has a recursion of two.

specular direction only.

This recursion concept is illustrated in Figure 18. Three initial hit points in three different regions under the coffee table are shown. The initial hit point (ray 1) in Figure 18a sends out shadow feelers to the sources to determine what their contributions are. Only one source can be seen from this location, the other is fully occluded. We then send out a second ray (recursion of two or ray 2) in the specular direction to see what the contribution is there. Again, this specular hit point sends out shadow feelers only to find out that it is fully occluded from both sources. Therefore the illumination from this location is due to one source only with no contribution from the specular bounce.

The middle-hit point (Figure 18b) finds out that both initial and specular shadow feelers are fully occluded from the sources. In this region, the radiance will be zero. Lastly, the initial hit point (ray 1) in Figure 18c traces shadow feelers to both sources and sees that they are both occluded. This time, however, the specular trace (ray 2) yields an energy contribution due to the fact that it can see both sources. As a result, the radiance for this initial hit point will be due to the energy contribution from the specular

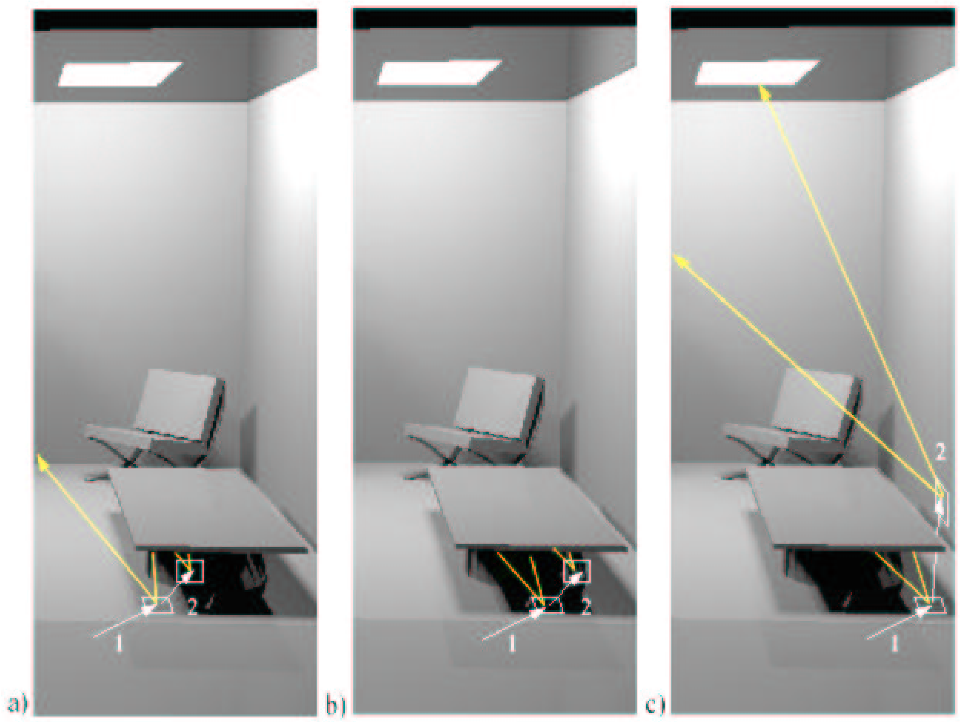

Figure 18: Diagram showing the important of recursion. Neither a) nor b) has an energy contribution from the specular bounce. Only c) propagates energy from the specular bounce to the initial hit point.

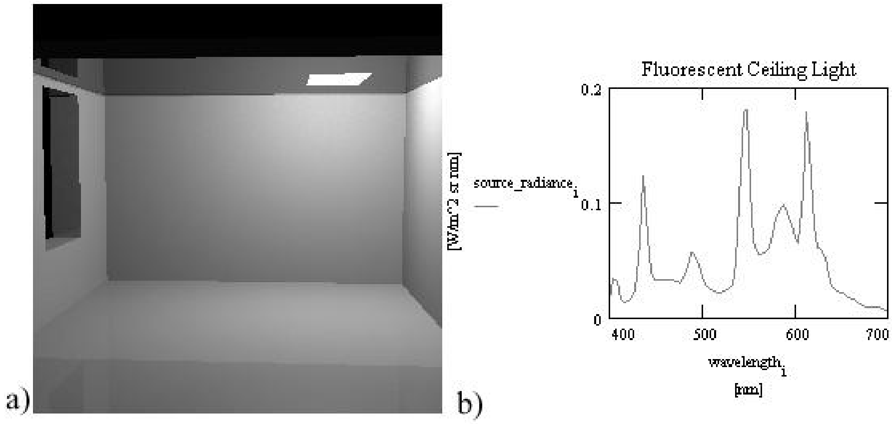

Figure 19: a) DIRSIG room with one area source. b) Spectral distribution of the area source.

bounce only.

It is clear from the results above that the DIRSIG model can fully simulate extended area sources to a (more than) reasonable degree. The big picture here, however, is that DIRSIG was not originally intended to simulate indoor scenarios. Nor was it thought that extended area sources were applicable. This is because most of the scenarios created using DIRSIG contain substantial altitudes or distances therefore making point source approximations valid. However, in support of requirements generated by high spatial resolution, low-light-level imaging scenarios, we have used a traditional ray tracer to generate radiosity type (or diffuse-diffuse) solutions. This type of photon/surface interaction becomes important when the user generates scenarios where point source approximations are no longer valid.

3.3.4 Quantitative Verification of Test Room Results

The following section validates the extended area algorithm and radiometry presented earlier. This is accomplished by first generating a DIRSIG room with one single extended area source in it (see Figure 19a). The source has a known spectral distribution (Figure 19b) and is read-in by the ray tracer.

Once the image has been rendered, individual pixel values can be "poked" resulting in radiance values as a function of pixel location. Therefore, we should be able to manually integrate the spectral source distribution and compare it to the result obtained by DIRSIG. Manual integration of the

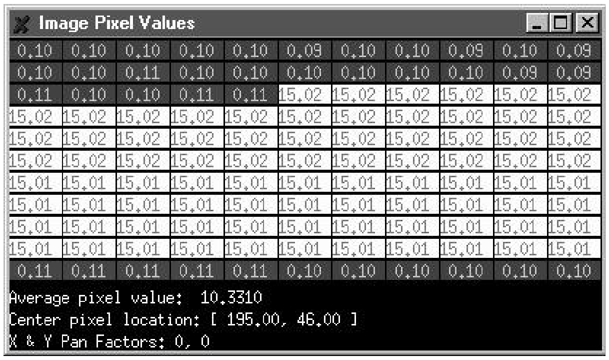

Figure 20: Section of area source in DIRSIG image displaying radiance values as a function of pixel location.

source distribution yields,

$$L = \sum_{j=0}^{N-1} \texttt{source\_radiance}_j \, \Delta\lambda \tag{9}$$

$$L = 15.058 \quad \left[\frac{W}{m^2 sr} \right]$$

We now choose a section of the light source in the DIRSIG image and examine the radiance values, so as to compare them to the manual integrated value obtained above. From Figure 20 it is evident that DIRSIG has integrated the source distribution correctly.

The next thing to calculate is the radiance value at a known distance from the source. This value will also be compared to DIRSIG's output. Since the geometry of the room is known, we can compute an irradiance directly below the source as,

$$\begin{array}{llll} \tau = 1 & R = 0.90 & xdim = 4ft & r = 9ft \\ \alpha = 0 & m = 1 & ydim = 2ft & \end{array}$$

$$E = r^2 \cos^m \alpha \int_{-\frac{xdim}{2}}^{+\frac{xdim}{2}} \int_{-\frac{ydim}{2}}^{+\frac{ydim}{2}} \frac{L\,\tau}{(r^2 + x^2 + y^2)^2} R \, dx \, dy \tag{10}$$

$$Lgnd = \frac{E}{\pi} \tag{11}$$

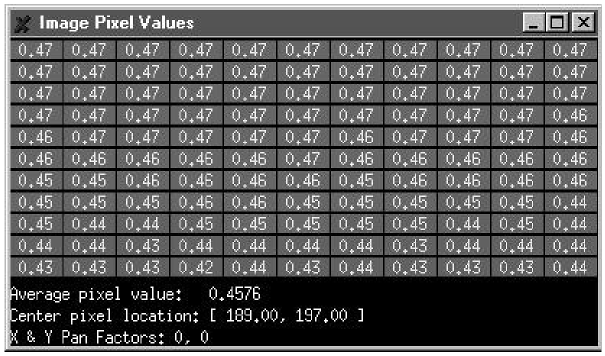

Image Pixel Values

0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47
0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47
0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47
0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.46
0.46	0.47	0.47	0.47	0.47	0.47	0.46	0.47	0.47	0.47	0.46
0.46	0.46	0.46	0.46	0.46	0.47	0.46	0.46	0.46	0.46	0.46
0.45	0.45	0.46	0.46	0.46	0.46	0.45	0.46	0.46	0.46	0.46
0.45	0.45	0.45	0.46	0.46	0.45	0.45	0.45	0.45	0.45	0.44
0.45	0.44	0.44	0.45	0.45	0.45	0.45	0.45	0.44	0.45	0.44
0.44	0.44	0.43	0.44	0.44	0.44	0.44	0.43	0.44	0.44	0.44
0.43	0.43	0.43	0.42	0.44	0.43	0.44	0.43	0.43	0.43	0.44

Average pixel value: 0.4576
Center pixel location: [189.00, 197.00]
X & Y Pan Factors: 0, 0

Figure 21: Section of ground directly below the source displaying radiance values as a function of pixel location.

$$Lgnd = 0.410 \quad \left[\frac{W}{m^2 sr}\right]$$

where r is the distance from the source to the ground, $xdim$ and $ydim$ are the dimensions of the source, τ is the transmission, R is the reflectivity of the floor, while α and m are the angle and shape value described earlier.

Similarly we examine a region of the DIRSIG image directly below the area source (Figure 21) and compare the ray-traced values to the one obtained by manual integration. Again, we see the values the ray-tracer generated are extremely close to what was obtained manually. This further solidifies DIRSIG's execution of the algorithm and radiometry presented earlier.

3.3.5 Extended Area Sources on Vehicle

We have just seen how area sources can be used in typical office lighting situations. We now extend this further by placing an area source(s) on a vehicle, such as the one illustrated in Figure 22 . Here we see two area sources used to simulated headlights on a vehicle. These headlights or source facets are used in the same manor as those found in the office scene earlier. The only difference is that the headlight facets have a source shape value of 20 while the office room light had a shape value of zero. The direction of the light energy follows that of the facet normal, which points slightly down and inward. This control of the direction and shape yields a more realistic looking vehicle headlight.

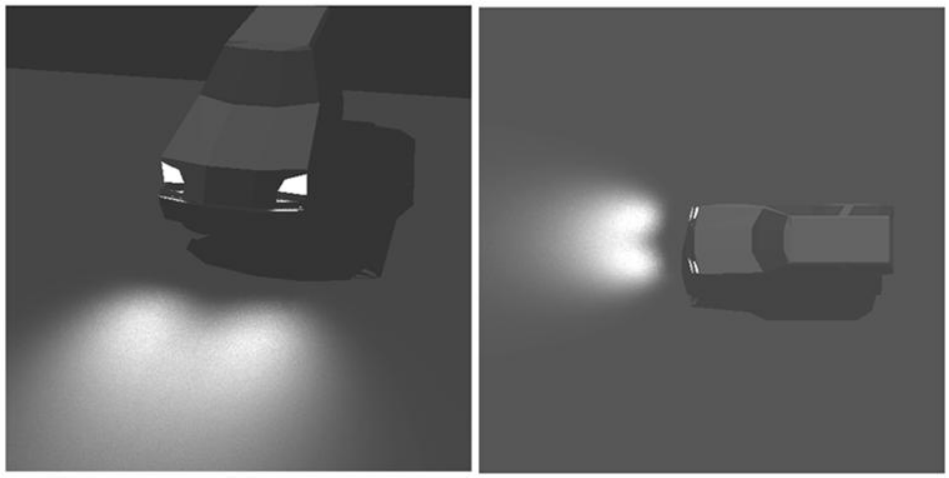

Figure 22: Vehicle with simulated headlights.

3.3.6 Summary of User-Controllable Variables

The last part of this section deals with the effect of varying some of the user-controllable variables discussed earlier. It is very important to note, however, that the proper use and understanding of these variables could potentially save the user significant amounts of run time. We will vary some of these parameters using the vehicle from 22. A brief discussion of each variable, including default values, follows with an illustrative example where applicable.

`SAMPLE_AREA_DENSITY` If a source has an area of A,we will compute a set of random sample points, M, within the source. A random subset of these points, N, will be used each time the source is checked. That way at rendering time, we don't have to compute the points to trace to, we just *randomly* pick some out of the pre-computed set (which should be larger than the number of rays we expect to throw at it). The number of points to make is the computed area, A, divided by the `SAMPLE_AREA_DENSITY`. The default value is 1 $[1/scene\,unit^2]$.

`MAX_SAMPLE_POINTS` This can be thought of as the ceiling or limiting value for the *sample point* value, M. If an area source is close and appears quite large, then the *sample point* value M may be large as well. The `MAX_SAMPLE_POINTS` value can limit the number *sample points* generated.

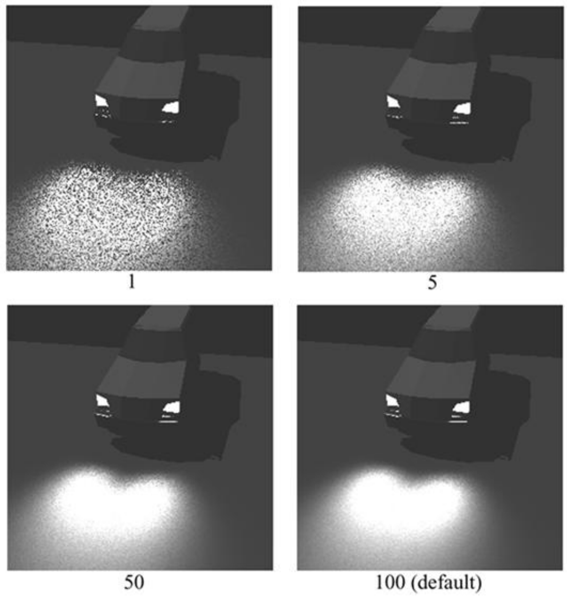

Figure 23: Effect of varying MAX SOURCE SAMPLES.

The default is 500 [*unitless*].

`MAX_SOURCE_SAMPLES` The value for the *source sample*, N, will be different at each hit point because the solid angle varies as a function of hit point location. Using the `MAX_SOURCE_SAMPLES` variable can put a limit on the value of N. If the hit point is very close, it doesn't take 1000 samples to determine the contribution. The default is 100 [*unitless*]. Some examples are shown in Figure 23 using values from Table 1. We simply vary the value from 1 to 100. At a value of 5, for example, this says that after calculating the solid angle *at a given hit point*, shoot no more than 5 rays at the source to assess the occlusion and irradiance.

Table 1: Values used for MAX SOURCE SAMPLES example.

`SAMPLE_AREA_DENSITY`	=1	$[1/scene\,unit^2]$	(default)
`MAX_SAMPLE_POINTS`	=500	[*unitless*]	(default)
`SOLID_ANGULAR_THRESHOLD`	=1E-6	[*steradians*]	(default)
`SOLID_ANGULAR_DENSITY`	=1E+4	$[1/scene\,unit^2]$	

`SOLID_ANGULAR_THRESHOLD` This is the minimum solid angle of a source

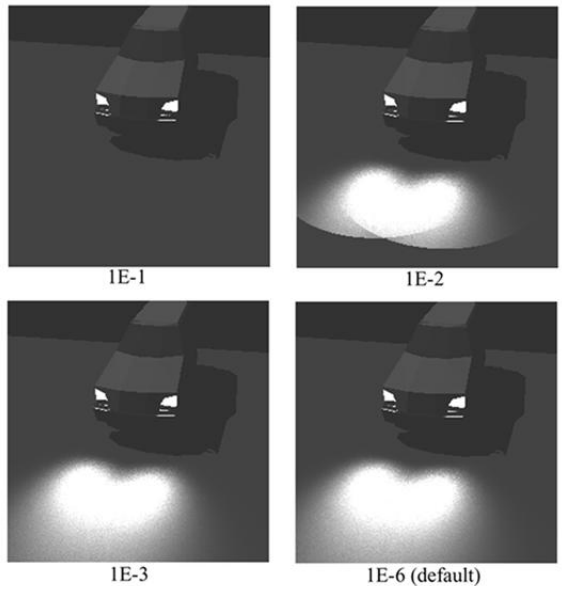

Figure 24: Effect of varying SOLID ANGULAR THRESHOLD.

to include in the calculation. If the computed solid angle of a source from a given hit point is *lower* than this, we ignore it and move on to the next source. The default value is 1E-6 [*steradians*]. Some examples are shown in Figure 24 using values from Table 2. We start off with a large value (1E-1) in which the source is ignored. At 1E-2 a threshold barrier is created in which certain hit points ignore the source. At 1E-3 smaller, all the hit points that can see the source, trace to it.

Table 2: Values used for SOLID ANGULAR THRESHOLD example.

`SAMPLE_AREA_DENSITY`	=1	[1/*scene unit*^{2}]	(default)
`MAX_SAMPLE_POINTS`	=500	[*unitless*]	(default)
`MAX_SOURCE_SAMPLES`	=100	[*unitless*]	(default)
`SOLID_ANGULAR_DENSITY`	=1E+4	[1/*scene unit*^{2}]	

`SOLID_ANGULAR_DENSITY` At each hit point we compute the solid angle, *Omega*, of the source. To figure out how many rays (or *source samples*, N) to throw, we multiply the apparent solid angle of the source by the `SOLID_ANGULAR_DENSITY`. The default value is 1E+3 [1/*steradians*]. Some

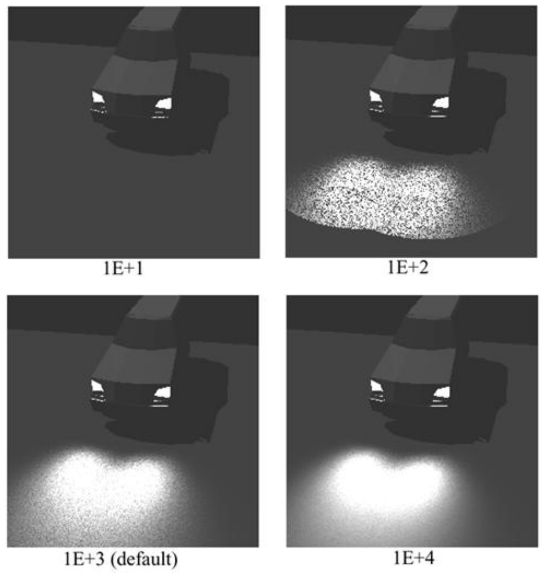

Figure 25: Effect of varying SOLID ANGULAR DENSITY.

examples are shown in Figure 25 using values from Table 3.

Table 3: Values used for SOLID ANGULAR DENSITY example.

`SAMPLE_AREA_DENSITY`	=1	$[1/scene\, unit^2]$	(default)
`MAX_SAMPLE_POINTS`	=500	$[unitless]$	(default)
`MAX_SOURCE_SAMPLES`	=100	$[unitless]$	(default)
`SOLID_ANGULAR_THRESHOLD`	=1E-6	$[steradians]$	(default)

`IRRADIANCE_THRESHOLD` Prior to actually tracing the ray to the point source, the logic checks to see if the energy contribution from the source will be significant. The integrated radiance for the current sensor bandpasses is precomputed when the spectral radiance curve is loaded for each source. During the rendering phase, the estimated integrated irradiance reaching the current location is computed (assuming no obscurations along the path to the source) and compared to a user controllable threshold. If the source does not contribute more than the indicated threshold, then the source is ignored and the computational burden of ray tracing and radiation propagation along the path is avoided.

This approach is favored in comparison to a simple distance threshold since it accounts for the spectral output of the source. The irradiance thresh-

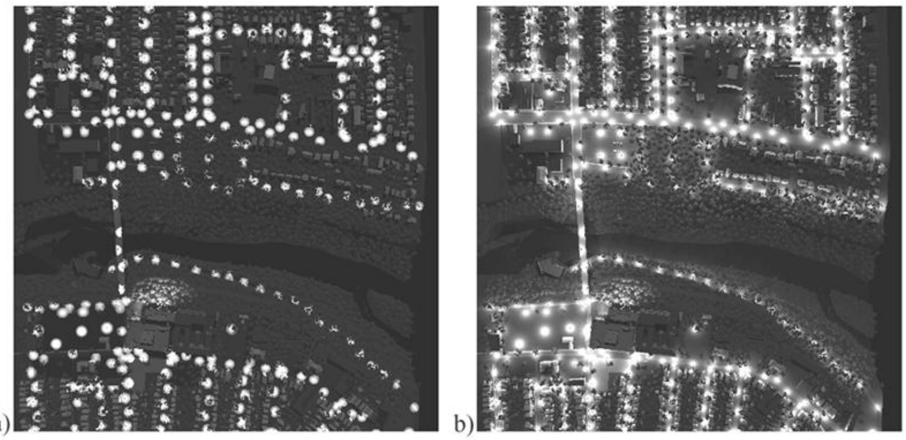

Figure 26: Effect of varying IRRADIANCE THRESHOLD. a) Is set to 1E-09 while b) is 1E-12.

old can be tuned to some desired level of fidelity by the user. If the threshold is set too high, the illumination field will appear to be "cut-off" (Figure 26). If the threshold is very low, the simulation will simply take longer. The ideal threshold for any given simulation should be about the magnitude of the direct and scattered energy from the major source (the Sun or the Moon). In future versions, the default threshold may be set to estimates of these major contributions. The default value is 1E-10 $[W/m^2]$.

3.4 Additional Implementation: Foxbat and Hawkeye Scenes

Once the advantage of shaped sources was seen, they were then implemented on larger real word scenes to further illustrate their utility. The first scene to incorporate a variety of shaped sources was the Foxbat scene (see Figure 27). Here, a series of point and non-point sources were used to simulate hanger lights, spotlight, and vehicle headlights. It is noticed that the geometrical shadowing on the ground and building looks more natural than would have been created with point sources alone. The spectral distribution of the sources used in this scene resembled that of a 2800 K planckian radiator (tungsten or tungsten-halogen).

Another scene that was populated with directionally shaped sources was the Kodak Hawkeye scene. A visible nadar view of this scene is shown in Figure 28. The source location information was obtained by performing an onsite evaluation of the surrounding areas.

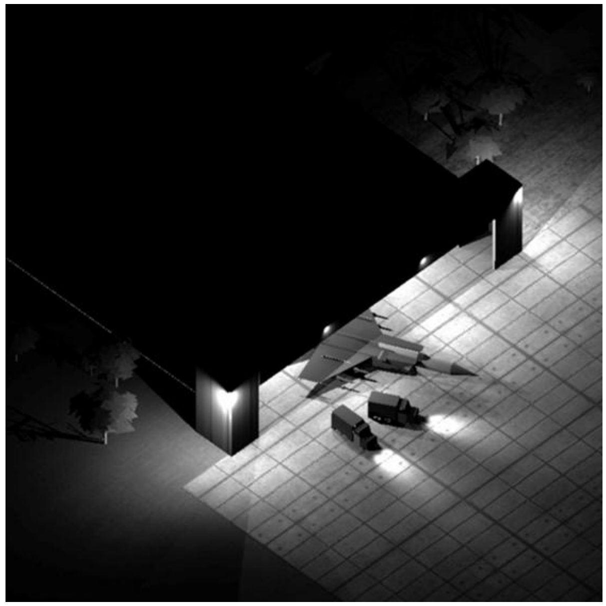

Figure 27: Foxbat scene with multiple directionally shaped point sources.

From the scene above, it was determined that there were two types of streetlights to be modeled. One of these lights can be seen in Figure 29. The first type is a standard streetlight found illuminating roadways and parking areas. These typically contain sodium or mercury vapor sources. Another type of streetlight to be modeled was that found in a typical neighborhood, smaller in size. These streetlights typically illuminate residential areas and walkways. More often than not, this type of streetlight contains a (low pressure) sodium source. The spectral distribution of such sources was measured and can be found in section 6 on page 61. Additionally, two of the Hawkeye buildings were re-designed in a CAD environment to incorporate 2 x 4-foot ceiling panel lights. Figure 30 illustrates the geometric locations of some of the ceiling lights in the Hawkeye buildings.

Once the sources were incorporated into the Hawkeye scene, a series of views were rendered. For comparison purposes, some of the rendered views attempted to match those created in previous DIRSIG runs that exemplified different phenomenology (*e.g.*, thermal). Figure 31, for example, shows a thermal image as well as a LLL image rendered from approximately the same view angle. It is seen that the interaction of the shaped light source with the ground in the LLL image looks as expected from a typical nighttime scene

Figure 28: True color nadar view of Hawkeye scene.

as viewed from overhead. The source shape value used for the roadway and residential streetlights was zero (*i.e.*, point sources). Since the image is in color, the orange-yellow color of the low-pressure sodium lights stands out.

Clearly, it is seen that this type of rendering further enhances DIRSIG's capabilities for generating hypothetical low-light-level scenarios. Additionally, since DIRISG has nighttime sources of illumination, it can now render a scene in any bandpass from 0.3 to 20 μm, day *or* night.

A nadar view of the Hawkeye scene was also rendered. This can be seen in Figure 32 along with a visible rendering of the same scene from the same view angle. Another view that was rendered contained the rose garden located near the Hawkeye building. This image is seen in Figure 33 along with its thermal equivalent.

A series of views were rendered to illustrate not only streetlights, but also shaped and non-shaped area sources on vehicles and in offices. These images can be seen in Figure 34. Here we see that the shaped headlights look fairly natural. The results would have been much different, and unnatural, had we used point sources. The building also has offices with room lights on. Notice how one can see some of the geometry *inside* the office. Again, although these images look ascetically pleasing, the focus is in DIRSIG's ability to accurately model the radiometry given an arbitrary geometric scene.

Figure 29: Large streetlight used in Hawkeye scene.

Figure 30: CAD drawing of Hawkeye building showing geometric locations of ceiling lights.

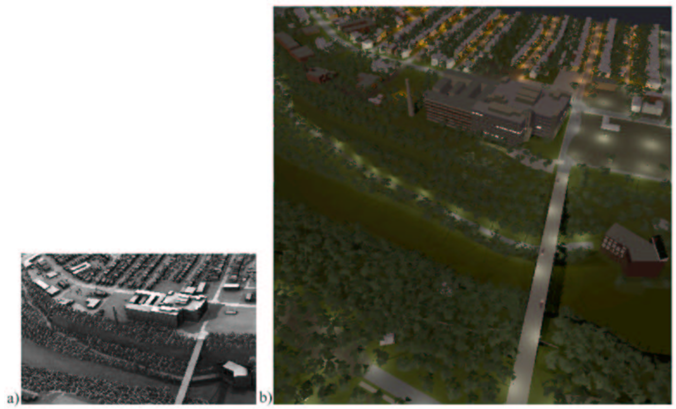

Figure 31: a) Thermal and b) LLL DIRSIG image of Hawkeye scene.

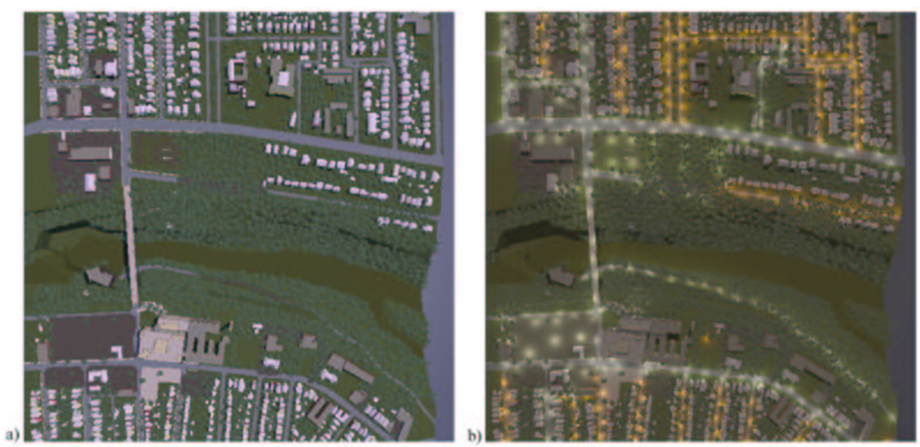

Figure 32: a) Daytime VIS nadar and b) nighttime LLL image of Hawkeye scene.

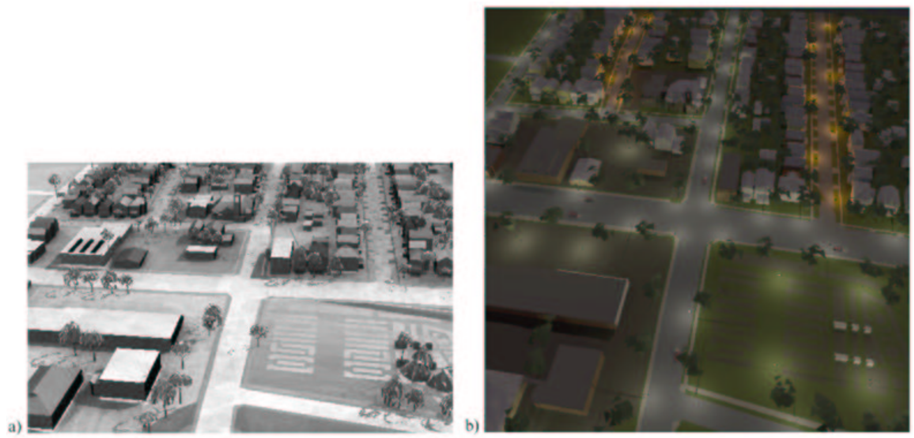

Figure 33: a) Thermal and b) nighttime LLL image of rose garden portion of Hawkeye scene.

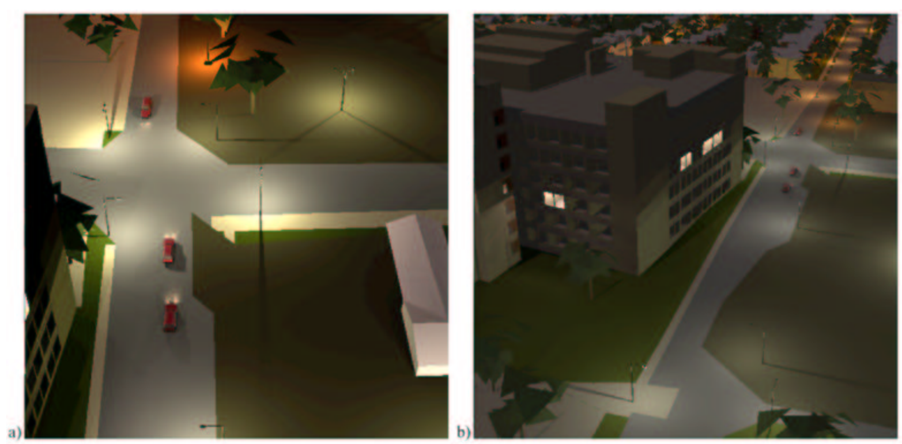

Figure 34: Images illustrating directionally shaped and extended area sources on a) vehicles and in b) offices.

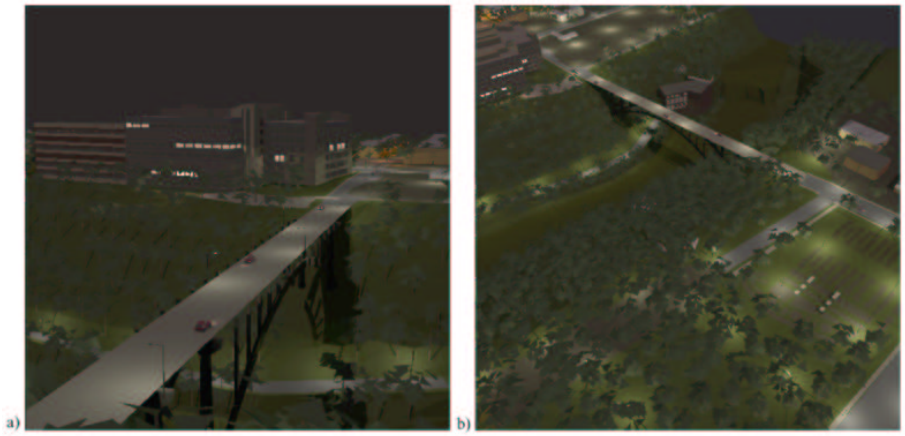

Figure 35: LLL images of a) Driving Park Bridge and b) rose garden.

From this, it is evident that directionally shaped sources have a clear distinct advantage over simple isotropic point sources. This type of flexibility in shaping and pointing the source dramatically improves DIRSIG's overall sense of realism, as was illustrated by modeling headlights and streetlights in a large urban scene.

A few more images of the Hawkeye scene were rendered from various viewpoints. These can be seen in Figures 35 and 36. Figure 35a shows the Hawkeye building with office lights on and the Driving Park bridge illuminated under a handful of streetlights while Figure 35b illustrates the interaction between streetlights and the tree canopy found in the rose garden. Figure 36 is another view of a typical residential area north of the Hawkeye facility.

4 Addition of Lunar Scattering

Prior to this report `make_adb`, the radiance database builder for DIRSIG, did not take into account lunar downwelled scattering. It was determined, through off-line runs of MODTRAN, that the downwelled term was indeed significant relative to the background spectral starlight distribution. Therefore a conscious effort was put in place to incorporate the lunar downwelled term in the atmospheric database. The results of this can be seen in Figure 37. Here we see spectral irradiance curves for various phase fractions of the moon. As expected, the irradiance decreases as the fraction gets smaller. The last curve in the data set is the spectral starlight distribution. It can

Figure 36: LLL image of Hawkeye scene.

be seen that the downwelled term for a 35% phase fraction source is very close in magnitude to this (starlight) distribution. However, one also needs to take into account the location (zenith and azimuth) of the source in the hemisphere above. The downwelled term can be different if the extraterrestrial source is located at directly above as opposed to being at the horizon, for a given lunar phase fraction. This effect was minimized by finding fractions that were located at similar zenith's and azimuth's in the sky dome (see Table 4). The downwelled runs were based on a visibility of 23 km.

Table 4: Image conditions for downwelled test cases.

	GMT Date	GMT Time	Local Time	Moon Elevation Deg	Moon Azimuth Deg	Moon Fraction %	Sun Elevation Deg	Sun Azimuth Deg
Full moon	9-6-98	0300	11pmEST	31.10	146.81	97.32	-32.33	321.15
3/4 moon	9-9-98	0600	2amEST	46.22	143.46	78.5	-40.04	17.04
1/2 moon	9-12-98	0900	5amEST	57.76	139.65	55.1	-19.413	64.10
1/4 moon	11-13-98	1000	6amEST	35.33	119.69	34.9	-22.02	94.39
New moon	9-22-98	0500	1amEST	-47.31	334.18	9.36	-46.43	358.79

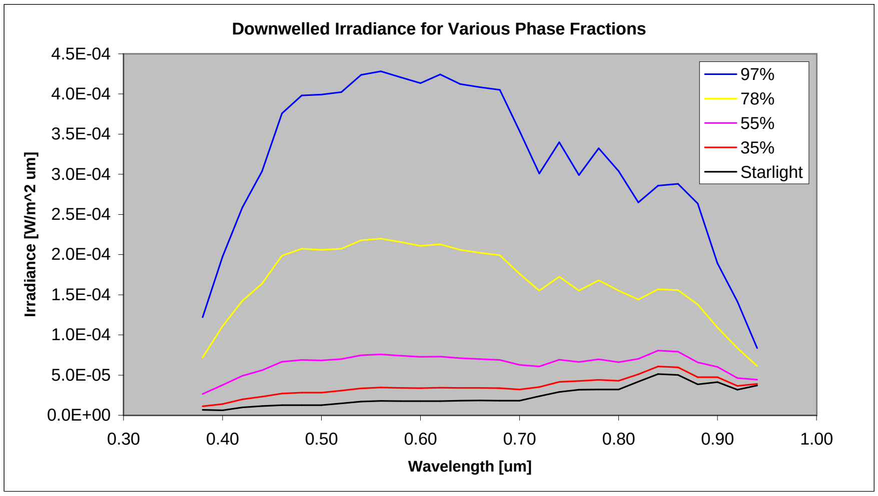

Figure 37: Downwelled irradiance for various phase fractions of the moon.

5 Low-Light-Level Validation Proposal

5.1 Introduction

In recent years, RIT's DIRSIG model has undergone visible and infrared validations by White and Kraska [1, 2]. This analysis focused on DIRSIG's ability to accurately predict radiometry for both the visible and thermal bands, as compared to truth data. More recently, however, upgrades to the SIG model have called for a new validation study. The latest version of DIRSIG includes man-made secondary sources and exoatmospheric sources such as the moon and starlight. These new enhancements enable the user to generate a wide variety of low-light-level scenarios. In general, the validation is crucial because the utility of the synthetic images is diminished if the output does not closely imitate the real world. As a result, the output from DIRSIG must be evaluated and assessed according to criteria such as spectral and radiometric accuracy, geometric fidelity, robustness of application, and speed of image generation.

A qualitative validation was performed by Ientilucci [3] that simply compared truth imagery to synthetic imagery based on geometry and crude laboratory parameters. These results were encouraging, but lacked quantitative value. This report, proposes a procedure that rigorously evaluates the low-light-level integrity of DIRSIG.

5.2 Validation Overview

In general, a SIG validation consists of 4 steps. 1) Collection of ground truth imagery, 2) generation of a synthetic equivalent of the truth, 3) calibration of instrumentation, and 4) development of metrics to quantitatively compare the data sets. For a low-light-level analysis, accurate ground truth collection is undoubtedly the most difficult to attain. This is because of the nature of the surround (controlled skylight) at the time of collection and the high sensitivity of the instrumentation used. Additionally, it should be noted that the output of DIRSIG is a per-pixel integrated radiance field image. Prior to this integration, the model takes into account *spectral* parameters such as sensor response, material reflectivity, and light source distribution. Therefore, it very important that these parameters be characterized properly before making any valid conclusions.

The following sections provide an inside look on how to create an outdoor (or indoor) test scene, what objects to include in the scene and what types of instruments should be used to image it. Additionally two methods of calibration are explained as well as two techniques for quantitatively evaluating the results. Lastly, sections are provided that deal with indoor calibration, simple error checking, and MTF calibration procedures.

5.3 Collection of Ground Truth

Real imagery should be obtained before any synthetic scenes are to be constructed, because it is somewhat easier to generate a synthetic scene based on a real image than it is to compose a real-life scene based on a previously built synthetic scene.

5.3.1 Illumination Conditions

Before the contents of the scenes are described, we make note of another important factor: *scene condition*. This pertains to the illumination conditions at the time of the collection. Perhaps the most difficult thing to control is stray light entering the scene from distant city lights, for example. On a previous collection by Ientilucci [3], all the sources of illumination for a large section of the campus (RIT) were turned off. To the naked eye this seem more than sufficient. However, a LLL sensor was still able to pick out distant sources of illumination. The solution here is to pick a location that is as remote as possible. Control of illumination is key for this is what we are trying to characterize.

5.3.2 Scene Construction

DIRSIG has the ability to model starlight, man-made sources, and varying phases of the moon. Therefore it would be a good idea to image scenes under all of these conditions. Six ideal scenes are proposed below that describe various combinations of starlight, moonlight, and light from man-made sources. These illumination conditions are summarized in Table 5.

Scene One. Here we image a scene that is illuminated by a full moon only. Actually, the phase fraction need only be close to 100%. This is because the SIG model can predict ephemerides for all the planets given parameters such as the date, time of day, latitude, and longitude. This scene lacks any man-made sources.

Scene Two. This scene is illuminated by the moon with a phase fraction of 0.5, or half moon. This scene also lacks man-made sources.

Scene Three. This scene is imaged under starlight conditions only. A key factor in this image is the lack of man-made source *and* moonlight.

The next three scenes are the same as the last three with the exception of the inclusion of man-made sources (the type of which is to be described later).

Scene Four. This scene contains a full moon and the presence of a man-made source(s).

Scene Five. Here we have a half moon and a man-made source(s).

Scene Six. Finally, we have a new moon condition with a man-made source(s). This scene aids in the SIG validation of secondary sources.

Common to all these cases is the assumption that the night sky is relatively free from cloud cover. Clouds will inherently complicate the validation process for they weigh in as yet another complex variable to be modeled.

It should also be noted that the man-made source portion of the validation could take place in a controlled indoor environment such as a laboratory (see section 5.7). In order for this to be a reality, DIRSIG would have to be able to run without an atmospheric database to reference (since there is no starlight or moonlight indoors). As of this writing, however, DIRSIG does not have this capability. For now we will image the sources outdoors.

Table 5: Summary of imaged low-light acquisition conditions.

	Full Moon	Half Moon	New Moon (starlight)	Man-Made Sources
Scene 1	x			
Scene 2		x		
Scene 3			x	
Scene 4	x			x
Scene 5		x		x
Scene 6			x	x

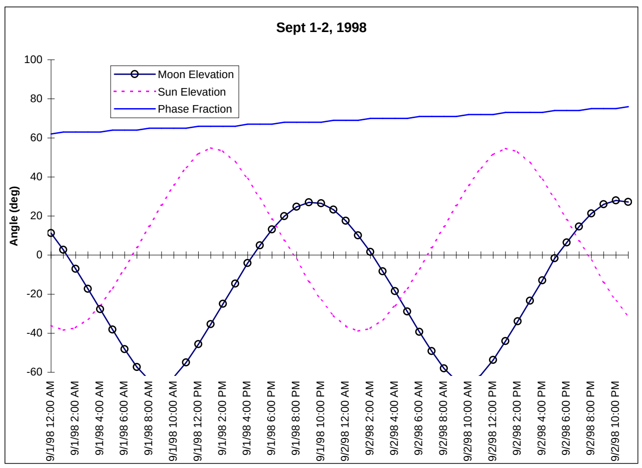

Figure 38: Output of ephemeris routine showing optimal data collection times for August.

5.3.3 Use of Ephemeris Utility to Aid in Collection

The use of an ephemeris utility can prove to be very useful when trying to predict optimum collection times. By using such a utility, one can predict which days will have a full moon, half moon, or new moon. One such program is call "ephem" and is provided by the National Optical Astronomy Observatories (NOAO). This program was previously used to predict the moon's position in the sky for a given time and day of year (see Figure 38). Ephem is an interactive astronomical program that displays ephemerides for all the planets. It is based on standards set forth by NOAO. The use of such a program could be invaluable in predicting which days are better candidates than others for image capture. Of course, it does not predict weather conditions so this will always be a factor to consider at the time of acquisition

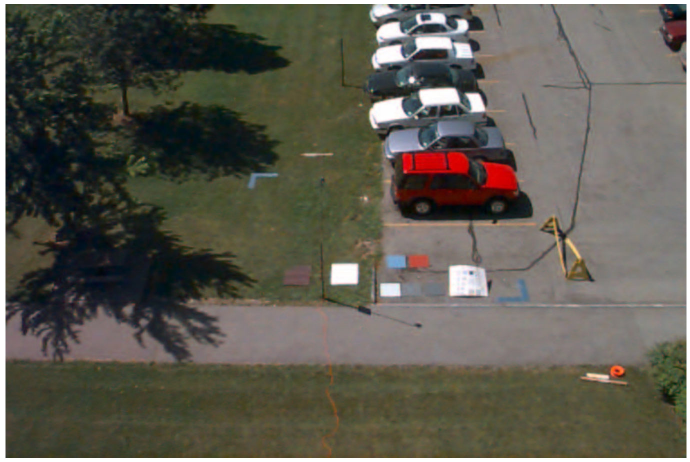

Figure 39: View of example layout for ground truth collection.

5.3.4 Ground Truth Scene Content

Once a remote site is determined, the next logical step is to construct a test scene that contains relevant targets, objects, and possible man-made sources. Suggested targets and objects are summarized in Table 6 with an example layout illustrated in Figure 39.

Table 6: Summary of targets and objects in scene.

Targets	Secondary Sources and Objects
Black Target	Sodium Vapor Source
White target	Mercury Vapor Source
Blue target	Tungsten or Tungsten-Halogen Source
Red target	
Light gray target	Human
Medium gray target	Vehicle
ark gray target	
Resolution target	
Aluminum targets	

Including plenty of solid colored targets in the test scene is a very important step in the validation process. This is because many ground control points (GCP) will be needed in order to get meaningful results from such metrics as rank order correlation (ROC) or RMS (see section 5.6). The GCP's will be solid areas that are uniformly illuminated such as the gray panels, colored panels, sidewalk, asphalt, and truck hood.

A USAF resolution target is always good to have in any collection when potentially evaluation the resolution of a system. The main thrust of this validation is not on the camera system but on DIRSIG's ability to accurately produce spectral radiance values. However, resolution characterization will play an important role if one is including the sensors MTF as part of the validation. This is discussed further in section 5.9. Most people, however, implement the sensor model as a post process operation. Here we propose the use a standard 1951 United States Air Force chart, which was designed for checking lenses used in aerial photography.

For visual purposes we have included objects such as a human and/or vehicle. Many low-light-level scenes seem to involve the ability to discern a truck, car, tank, or human. Therefore, it may be beneficial to include such objects. It should be noted, however, that the inclusion of such objects is more for qualitative purposes. That is, they do not play an important role in the quantitative aspect of the validation.

5.3.5 Equipment Used in Collection

This section details the equipment that might be used in a collection such as the one stated above. The first thing to be measured should be the properties of the targets and objects in the scene. Current users of DIRSIG can reference a standard set of emissivity curves for generic materials, however, some of these curves are not entirely correct in the visible nor are they even available. It is therefore necessary to develop a set of emissivity/reflectivity curves specific to objects in the scene.

Spectrometer

The radiance of objects can be measured with a variety of spectral radiometers such as the Spectra Colorimeter, model PR-650 by Photo Research or the one manufactured by Analytical Spectral Devices (ASD). The latter has a spectral range form 350 to 2500 nm in 1 nm increments while the former has a range of 380 - 780 nm in 4 nm increments. The reflectance can be determined by first measuring the radiance level from a piece of pressed polytetrafluoroethylene (PTFE) that is placed in the area of interest. Then by removing the disk, one measures the sample area underneath. The reflectance is found by dividing the sample spectra by the reference spectra. This is a valid approach because the pressed PTFE exhibits a remarkably high diffuse reflectance over the range of 200 - 2500 nm [4]. Its reflectance is 99% or higher over the spectral range of 350 - 1800 nm. Finally, the spectral

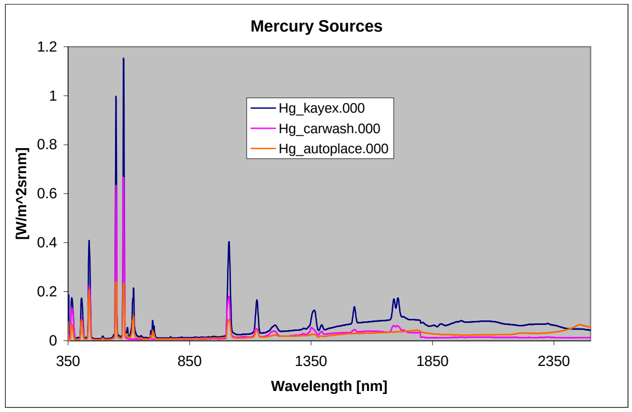

Figure 40: Spectral distribution of mercury sources.

emissivity curves are computed since,

$$\epsilon(\lambda) = 1 - r(\lambda) \tag{12}$$

The source distributions can also be measured with a radiometer such as the ASD. An example of a source distribution as measured with an ASD is illustrated in Figure 40.

Sensor System

We now turn to the camera system that will measure the radiance field from the scene. A common sensor that is used for low-light-level imaging is an *intensified CCD* (ICCD). This design uses a photocathode to convert photon to photoelectrons. It then uses a micro channel plate (MCP) to amplify the photoelectrons. Finally the photoelectrons impinge themselves onto a phosphor so as to be read by a CCD. These systems are inherently noisy and usually have poor MTF response. This degradation is mainly due to the photon passing through multiple conversion stages as well as the amplification process. Another type of low-light-level device is an *electron bombarded CCD* (EBCCD). These systems accelerate a photoelectron onto a back-thinned CCD directly therefore omitting multiple conversions while improving overall resolution. The difficulty with such systems lies in the fabrication of the thinned CCD. If one performs the source part of the collection indoors, see section 5.7, then a standard CCD system will suffice. Here all the gain steps found in the LLL device are eliminated, thus making calibration much simpler and perhaps more accurate.

Imaging RGB

We have just discussed imaging the scene across the entire bandpass of the sensor. At the same time, one could insert a series of narrow bandpass filters so as to perform a spectral validation as well. For this LLL application the bandpass filters may be quite large, seeing how we are photon starved to begin with. Therefore, one could capture four images of a scene. The first being a very broad band image (the photocathodes response) and the remaining three as seen through a red, green, and blue filter.

Auto-Gain Control (AGC)

One of the biggest problems with commercially available (LLL) sensor systems is the inability to over ride the automatic gain control or AGC. The AGC prevents the camera system from overloading or burning out due to over amplification. However, this means of protection is usually a non-linear process therefore making it difficult to interpret imaged data. With a little skill and help from the manufacture, one can over ride the AGC thus yielding output voltages that are linearly related to input signal flux.

Regardless of the sensor, they all send out a voltage signal proportional to the incident light level. If the device reads the CCD out at 33 ms, then the output is in the form of a standard video signal, which is usually 1 Vpp. It is this voltage level that needs to be captured by a frame-grabbing device and ultimately calibrated.

Capture Device

The frame grabber will digitize one frame of the signal imaged. It is this image that will be used as a means of comparison to the still-frame radiance field imagery generated by SIG model. In a previous collect, a simple frame grabber was used that had a dynamic range of 8 bits with a resolution of 640 x 480. This proved to be inadequate while attempting a calibration. Therefore, it is recommended that a 12-bit frame grabber be used at the very least. This provides about 3 orders of magnitude in dynamic range. The intra-scene dynamic range may be quite large, however, possibly up to 5 orders of magnitude. This would call for a 16 or 17 bit digitizer. This is ultimately a function of the imaging device, however, so it is important to select an imaging device that has a wide dynamic range. In terms of spatial resolution, it was found that 640 x 480 generated images of poor quality. Difficulties arose when imaging a sine wave chart while performing an MTF characterization. Here it is recommended that the spatial resolution be 1024 x 768 (or higher) so as to minimize such difficulties.

5.4 SIG Scene Creation and Running the SIG Model

Once a truth scene has been built, a SIG equivalent of it must be constructed. It is these two scenes that will, ultimately, be compared in the validation.

Many of the DIRSIG scenes that are simulated originate from drawing packages such as AutoCAD or Rhinoceros. With these packages one can draw wire frame models of parts and objects. These parts then get assembled into larger scenes to be ray traced. It should be noted, however, that during the collection itself one needs to accurately log all the locations and dimensions of various targets and objects. This information is important in the recreation of the truth scene. A summary of the steps in the overall SIG process can be outlined as follows:

- Construction of individual parts in AutoCAD or other drafting program
- Assembly of individual parts into objects
- Assembly of objects into an entire scene
- Computing the polygon normal vectors
- Exporting the scene from AutoCAD into the DIRSIG Geometric Database
- Insertion of spectral emissivity (ems)
- Insertion of spectral sources (int)
- Insertion of material properties (mat)
- Insertion of weather database (wth)
- Building the atmospheric data base for the scene (adb)
- Defining the sensor characteristics (rsp)
- Running the DIRSIG model on the scene
- Extracting images from the output files

Key files to consider in this list are the emissivity, source, material, weather, sensor, and atmospheric database files. The data that is entered into these files comes from parameters acquired at the time of the collection. The spectral emissivities of objects and sources are measured by the techniques explained earlier. Material attributes, such as specific heat, thermal conductivity, mass density, and exposed area mainly play a role in thermal analyses. The weather file also is used extensively in thermal predictions, however, parameters such as temperature, pressure, humidity, dew point, wind, sky fraction, cloud cover, rain, and insulation can effect a low-light-level analysis. Therefore, these parameters should also be recorded 48 hours before the time of the collection, in one-hour increments. Sometimes weather information can be obtained from a source such as the National Oceanic and Atmospheric Administration (NOAA) after collection time. Creation of the atmospheric database is undoubtedly the most important parameter. This file contains all the information about the irradiance from lunar and skylight sources. Therefore, what is needed here is the actual date and time of the collection, ground altitude, latitude, longitude, atmospheric conditions (contained in a MODTRAN tape 5 file) and previously created weather file. Lastly, the sensor response has to be taken into account. DIRSIG reads in a file that contains a sampled version of the sensor response that will be used to image the ground truth scene.

5.5 Calibration of Instrumentation

Now that a scene has been created and the spectral attributes of various materials and sources measured, we need to calibrate the sensor system that will image the scene. This is important for we need a way to convert the digitized values (in digital counts) to radiance values. In other words, we need to get both images in the same space, whether it be digital counts or radiance. Here we will forgo DC space for radiance space. Once this is known, the gain and bias factors from the calibration can be applied to the digitized imagery so as to convert them back to radiance space for comparison to SIG radiance imagery later on.

5.5.1 Radiance Calibration

As mentioned before, the calibration is needed to covert the digitized values back to radiance space. Two methods for this calibration are presented. The first technique utilizes neutral density filters while the second alters the illumination level of the source by using an aperture. It is believed that the second method may be more robust since the method involving ND filters may alter the spectral character of the source not to mention the potential for transmission inaccuracies.

ND Filter Approach

The general layout is illustrated in 41. We start off with a stable light source that is connected to a regulated power supply. The current through

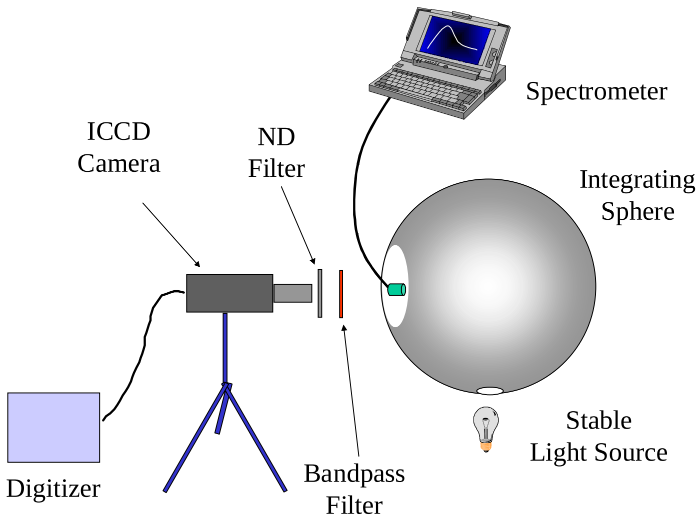

Figure 41: Calibration setup using integration sphere and ND filter.

the source must remain constant or else the total flux out of the source will change during the calibration. The source maybe of a tungsten or tungsten-halogen type with a very low flux output.

Once the source is selected, it is placed in an integrating sphere. The camera is placed at the exit port of the sphere so that light can uniformly fill the cameras field of view. It is this uniform field that will get digitized. The next step is to determine the integrated radiance that the camera will see. There are a couple of ways to do this. We could calculate the transfer of energy from the light source to the exit port using the following relationship

$$E_{out} = \frac{\Phi R}{4\pi r^2(1-R)} \tag{13}$$

where R is the integrating sphere surface reflectance and Φ is the radiometric flux. However, there is an easier technique that simply involves measuring the radiance at the exit port with a spectrometer. This technique puts the trust on the calibration of the spectrometer. Therefore, be sure the spectrometer has been recently calibrated. If we place a spectrometer, such as an ASD, at the exit port we can obtain the spectral distribution from the source and sphere together. We then simply multiply this result by the spectral transmission of the ND filter, bandpass filter, and camera response. To obtain the final input value that the sensor will see, we simply integrate over the bandpass. That is

$$L = \int L_{ASD}(\lambda)\beta(\lambda)\tau_{ND}(\lambda)\tau_{BP}(\lambda)\, d\lambda \tag{14}$$

where L_{ASD} is the radiance measured with a spectrometer, β is the normalized spectral response of the photocathode or camera system, τ_{ND} is the transmission of the ND filter, and τ_{BP} is the transmission of the bandpass filter. If the spectral response of the ND filter is unavailable, one could assume a flat response and measure the transmission density value (D_τ) with a densitometer. This can then be converted to a transmittance (τ_{ND}), which then modulates the integrated radiance value. That is

$$L = \tau_{ND} \int L_{ASD}(\lambda)\beta(\lambda)\tau_{BP}(\lambda)\, d\lambda \tag{15}$$

$$L = 10^{-D_\tau} \int L_{ASD}(\lambda)\beta(\lambda)\tau_{BP}(\lambda)\, d\lambda \tag{16}$$

It is this reduced radiance signal that will get digitized when the exit port radiance fills the entire field of view of the camera. Statistics about the digitized image, such as mean and standard deviation, can then be obtained. This calibration is performed for a series of camera gain settings thus yielding a look-up table that relates digital counts, radiance, and camera gain for each bandpass (*i.e*, red, green blue, VIS). The gain and bias parameters can then be applied to the digitized imagery to bring it back to radiance space. That is

$$L = g\, DC + b \tag{17}$$

where g is the gain and b is the bias. It should also be noted that the same f-number should be used as that used during the collection period.

Apertured Source Approach

The apertured source technique relies on cutting down the flux level by simply placing an aperture in front of the source instead of using ND filters in front of the camera (see Figure 42).

As mention earlier, ND filters add another (possibly unwanted) variable in the integration step. By simply placing an aperture in front of the source, we reduce the flux level with out altering its spectral character. The procedure and transfer equations are the same as before except for the omission of the transmittance term. That is

$$L = \int L_{ASD}(\lambda)\beta(\lambda)\tau_{BP}(\lambda)\, d\lambda \tag{18}$$

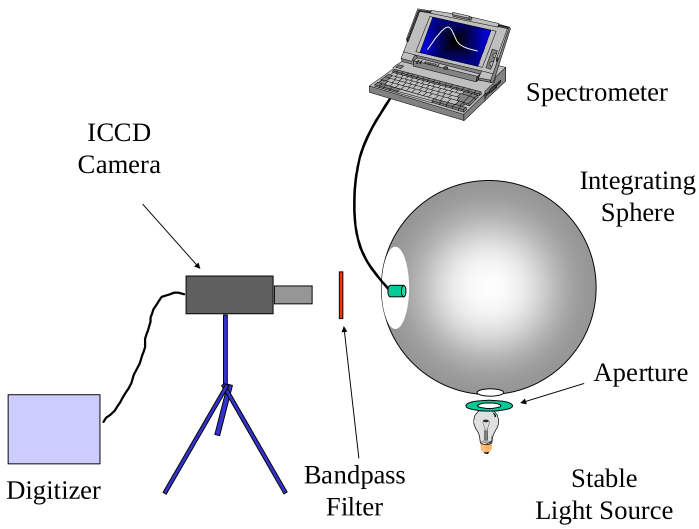

Figure 42: Calibration setup using integration sphere and apertured source.

Again, to obtain the radiance value we aperture the source, measure the exit port radiance with a spectrometer, perform the multiplication with the sensor response, and integrate over the bandpass. To populate the look up tables, this process is repeated numerous times as a function of camera gain. Finally the gain and bias parameters are applied to the digital count imagery to bring them into radiance space.

5.6 Quantitative Comparison

The final step in the validation process, is the comparison of ground truth and SIG imagery. This can be done after it has been determined that the two images are reasonably matched to one another. Several different methods of comparing images are available, including root mean square (RMS) error and rank order correlation (ROC). These techniques are described in the following sections.

5.6.1 Rank Order Correlation (ROC)

Rank order correlation is used to evaluate the relative contrast produced in a synthetic image as compared to truth. Relative contrast in an image is important for both human and computer based classifiers in evaluating an image and detecting specified objects within a scene. Each object of a specified number of objects in the scene are given a ranking in terms of the brightness of the object. This brightness ranking can then be used to compare a synthetic image with a real or truth image. If the contrast rankings in the synthetic image do not closely resemble the rankings of the truth image, a problem has been detected in the synthetic image generation process that will then result in a error when computer or human assessment of the synthetic image is performed.

Once the objects in both the synthetic and the real image have been ranked according to their brightness, or digital counts values, the comparison and, hence, assessment of the synthetic image can be accomplished. An initial evaluation can be seen by graphing the rank order in the truth image versus the rank order in the synthetic image. Ideally, this would be a perfectly linear graph with a one-to-one target correlation. Any difference in ranking between the DIRSIG image and the truth image is indicative of a contrast reversal for that object in the DIRSIG scene. By comparing images at different times of the day or night using the rank order correlation method, a good assessment of the synthetic image generation process can be found.

By examining the ROC of each corresponding object in the two scenes, an overall ROC coefficient can be assigned to the image. The Spearman rank order correlation is defined as:

$$\rho = 1 - \frac{6\sum(R_i - R_i')^2}{n^3 - n} \tag{19}$$

Where p is the correlation coefficient for each image pair at a particular time of day (TOD), n is the number of samples (or ground control points), R_i is the rank in the truth image for the i^{th} object, and R_i' is the rank in the synthetic image for the i^{th} object.

This overall rank order correlation can then be used in comparing the overall accuracy of the synthetic images. However, the one major problem with rank order correlation comparisons is that it is insensitive to overall radiometric gain or bias errors. To detect these gains and biases in the radiometry, the root mean square (RMS) error method of comparison is needed which is described in the next section.

5.6.2 Root Mean Square (RMS)

While the rank order correlation can help detect radiometric problems in individual objects, the root mean square (RMS) error method is useful in detecting overall problems in the synthetic image generation process. The root mean square error method helps to detect overall gain or bias problems in the radiometry of the synthetic image caused by errors in the atmospheric parameters of sensor characteristics. The actual RMS error is found by comparing the mean radiance values of objects in the synthetic image with the same objects in the truth or real image.

$$RMS = \sqrt{\frac{1}{n}\sum(L_i - L_i')^2} \tag{20}$$

Where n is the number of objects, L_i is the rank in the truth image for the i^{th} object, L_i' is the rank in the synthetic image for the i^{th} object.

5.6.3 Comparison Comments

Both the ROC and RMS techniques are applied to each image scenario (*i.e.*, full moon, starlight, etc.). Furthermore, each image scenario contains four band passes that need to be evaluated. This may produce anywhere from 4 to 12 sets of images to analyze and compare. The ROC evaluation uses the ranked values and will produce a correlation coefficient between zero and one. Additionally, the RMS metric can be performed in either radiance or digital count space.

5.7 Indoor Calibration of Sources

Unlike the moon and stars, man-made light sources can be set up in an indoor laboratory for evaluation. This may be beneficial in two ways. The first is the fact that a standard CCD could be used in place of the ICCD. The ICCD is used because the light levels are extremely low and using a source indoors may not require such sensitivity. The second relies on the fact that, indoors, the illumination conditions can be very controlled and isolated.

The procedure for the validation is exactly the same as what was stated earlier for the outdoor scenes. These steps are summarized below.

- Create indoor scene
- Image it with a regular CCD device (RGB and VIS)
- Build the CAD equivalent of the scene
- Run the SIG model to get the radiance field (RGB and VIS)
- Calibrate the CCD so as to convert DC values to radiance values

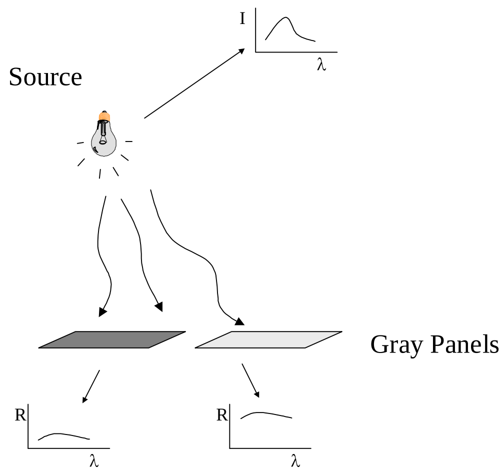

Figure 43: Example of indoor validation setup.

- Compare and contrast images using metrics such as ROC and RMS

5.8 Simple Sanity Check, No Imaging

This section deals with simple checks one can do that don't involve any imaging of the scene what so ever. These might be called "quick and dirty" checks. None the less, they provide meaningful results.

The simplest check involves the indoor scene just described. We first create the scene. An example of such a layout is illustrated in Figure 43.

Before the collection, the spectral character of the source and reflectance of the solid colored targets can be measured. During the collection one might also measure the radiance coming form the targets at specific locations. When it comes time to run the SIG model, all of these parameters are taken into account. Since there is no camera imaging the scene, the sensor response in the SIG model is set to one. The validation then takes the form of spectrally comparing the truth measurements to the SIG predicted measurements.

Illumination falloff errors are eliminated if the targets are uniformly illuminated. This will, more than likely not be that case, since the source and target are in such close proximity to one another. Therefore, we will have to take into account projected area effects.

The geometry for the radiometric calculations can be seen in Figure 44. It is first seen that the irradiance from the direct incident ray, E_o, is

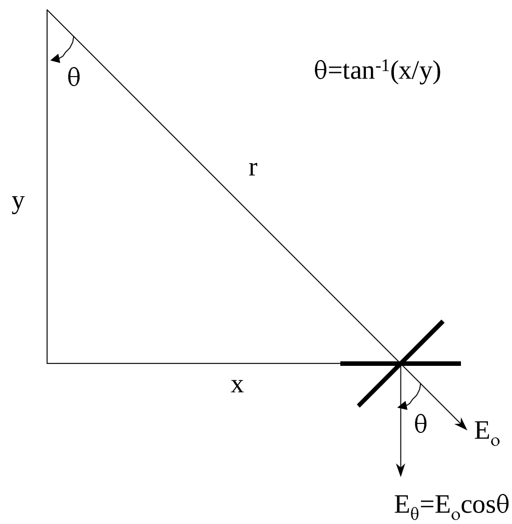

Figure 44: Geometric description of projected area effects (cos θ).

$$E_o = \frac{I}{r^2} \quad \left[\frac{W}{m^2}\right] \tag{21}$$

We then compute the irradiance onto a normal surface that is rotated through an angle θ, as illustrated in Figure 44.

$$E_\theta = E_o \cos\theta \quad \left[\frac{W}{m^2}\right] \tag{22}$$

Finally, we convert to radiance assuming approximate lambertion behavior.

$$L_\theta = \frac{E_\theta}{\pi} = \frac{E_0 \cos\theta}{\pi} = \frac{I \cos\theta}{\pi r^2} \quad \left[\frac{W}{m^2 sr}\right] \tag{23}$$

Once we know the spectral output of the source, reflectivity of the targets, and projected area effects, we can calculate the resulting radiance as seen by the sensor with a response of one. That is

$$L = \frac{\cos\theta}{\pi r^2} \int I(\lambda) R(\lambda)\, d\lambda \quad \left[\frac{W}{m^2 sr}\right] \tag{24}$$

where I is the intensity of the source, θ is the angle the source makes with the point of interest on the panel or ground, R is the reflectivity of the target, and r is the distance from the source to the point of interest on the panel or ground. This is the calculation that DIRSIG is doing therefore we are simply doing comparison to see if the SIG model performed this calculation. This type of "sanity checking" with out having to collect imagery can prove to be very useful and can be applied to other parts of the validation as well. These above steps are summarized below.

- Create indoor scene
- Measure spectral character of source and targets
- Create equivalent SIG scene
- Set sensor response to one and run SIG model to get integrated radiance field
- Calculate, based on measured spectra, what the integrated radiance should be at various points on targets
- Select the same points in SIG scene on targets and compare radiance values

5.9 MTF Calibration

As part of the calibration, one may desire to evaluate the *resolution* performance of the SIG imagery. In order to do this, the sensor system MTF would have to be calculated for comparison purposes. Once the sensor MTF is found, a limiting frequency value can be found at some specified value of the MTF. This limiting frequency can then be compared to the limiting frequency found on an USAF chart, for example, located in the SIG scene. The following sections describe how to read the resolution for and USAF chart and measure the MTF in voltage and digital count space.

5.9.1 Aerial Resolution

The spatial resolution of a lens is normally expressed in terms of line pairs per millimeter (lp/mm) in the *image plane*. We can calculate the resolution, at the *target*, by using the following equation

$$resolving\ power = 2^{m+\frac{n-1}{6}} \quad [lp/mm] \tag{25}$$

where m is the group number and n is the element number. By factoring in the magnification ($M = f/d$), we can calculate the resolution at the *image plane*. Usually one picks the smallest "block" where all horizontal and vertical spaces and bars are resolvable. From this, the values for m and n are obtained and one can then compute a value for the limiting resolution.

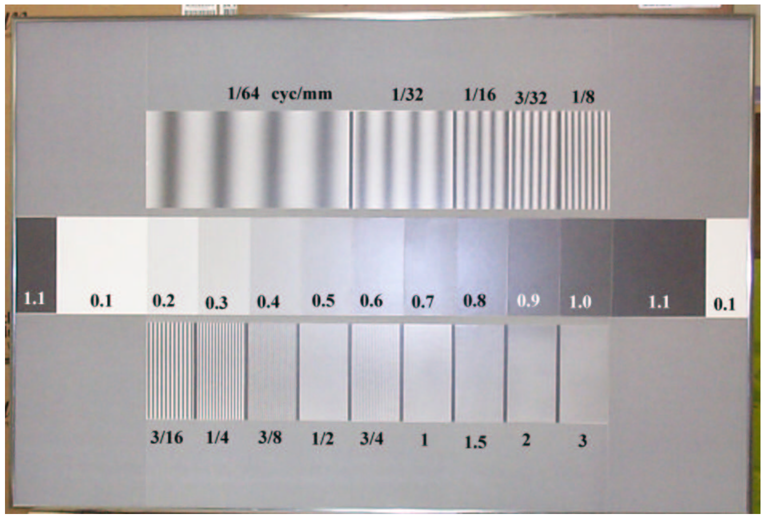

Figure 45: Typical sine wave target used to calculate sensor MTF.

5.9.2 Measurement of MTF in Voltage Space

The measurement of the voltage MTF pertains to accessing the "video" voltage signal that exits the sensor. A typical target for evaluating the MTF performance of an imaging system is illustrated in Figure 45. The chart is set up to have a series of varying sinusoidal patterns, each with a different frequency. At the center of the chart is a gray scale for calibration purposes.

The first thing that is done is to measure the reflection density (D_r) of each of the panels in the gray scale target. Once this is done the reflectance (R) can be computed by

$$R = 10^{-D_r} \tag{26}$$

The target is then imaged with the sensor so as to obtain voltage readings that correspond to each gray scale patch. In order to do this, one needs an oscilloscope capable of isolating one video scan line at a time. Once the voltages are measured, a calibration curve can be generated relating reflectance and voltage out. From this a calibration function of the following form can be generated

$$R = m\,V_{out} + b \tag{27}$$

This can be used to convert voltage to reflectance, which will be needed for the MTF calculation. Next we image the sine wave patterns on the target. The frequency, in cyc/mm, at the object plane (f_{obj}) is usually given for the chart. We simply factor in the magnification ($M = f/d$) to determine the frequency at the image plane (f_{img}). That is

$$f_{img} = M\, f_{obj} \tag{28}$$

Next, the sine wave patterns are imaged one at a time. An oscilloscope will show the voltage equivalent of such patterns. From this the minimum (V_{min}) and maximum (V_{max}) voltages are measured. These voltages are then converted to minimum (R_{min}) and maximum (R_{max}) reflectance's using the calibration equation above. Once this has been accomplished, the output modulation, for each pattern, can be found by

$$Modulation_{out} = \frac{R_{max} - R_{min}}{R_{max} + R_{min}} \tag{29}$$

The input modulation is found in a similar manor, except the reflectances are obtained by converting measured reflection densities found on the chart itself. Once the input and output modulation are found, the overall MTF is computed by

$$MTF = \frac{Modulation_{out}}{Modulation_{in}} \tag{30}$$

This process is repeated for each sine wave pattern or frequency. The MTF values are then plotted against their corresponding frequency values at the image plane (f_{img}).

In order to compare this to the USAF chart, a limiting frequency on the MTF plot must be established. A typical value is found by selecting the frequency found when the MTF is at 10%. It should be noted, however, that this is usually a fairly subjective decision. One last comment to be made is the fact that the resolution stated on the USAF chart is given in terms of line-pairs per mm (lp/mm) where as the modulation values were obtained from a chart that was in cycles per mm (cyc/mm). That is, one target is based on frequency patterns resembling square waves while the other is based on that of sine waves. Therefore, we are making an approximate comparison between the two. Conversion of this square-wave MTF to a sine wave MTF is accomplished by using the Fourier transform at a given frequency [5]:

$$T(f) = \frac{\pi}{4}\left[S(f) + \frac{S(3f)}{3} - \frac{S(5f)}{5} + \frac{S(7f)}{7} + \frac{S(11f)}{11} + \dots\right] \tag{31}$$

where $T(f)$ is the sine-wave MTF, $S(f)$ is the square-wave MTF, and f is the spatial frequency in $cycles/mm$. By calculating several values of $T(f)$ from the known square-wave function, the sine wave MTF can be determined and used for camera system optical image transfer analysis.

5.9.3 Measurement of MTF in Digital Count

Space The measurement of the MTF here pertains to including a digitizing device in the imaging chain. The process is very similar to the one stated above so we will omit some of the detail for brevity sake.

The first thing that is done here is to create a calibration curve like that obtained in the previous section. The reflection densities (D_r) are measured and converted to reflectance's (R) via the same procedure as before. We then digitize the gray scale and extract the corresponding digital count (DC) values. This generates a new calibration equation of the form

$$R = m\,DC + b \tag{32}$$

This equation is used to convert digital counts to reflectance's. Once this is established, each sine wave pattern on the chart is digitized and DC_{max} and DC_{min} values are obtained. These values are then convert to reflectance space to obtain R_{min} and R_{max}. The rest of the procedure is exactly as that stated earlier.

6 Spectral Distribution of Sources

6.1 Gas Discharge Sources

The spectra of a series of gas discharge sources were measured with a spectral radiometer from 350 to 2500 nm in 1 nm increments. The radiometer used was an ASD (Analytical Spectral Devices). This instrument was also cross-calibrated, in some cases, with a spectraphotometer (Photo Research PR650).

6.1.1 Sodium Lamps

Two high-pressure sodium lamps were measured. The first was a typical streetlight illuminating a busy roadway. The second was found on the side of a building illuminating a large parking area. A daylight picture of these sources can be seen in Figure 46. The corresponding spectral distributions can be seen in Figure 47 along with a blow up of the visible region.

Figure 46: a) Sodium streetlight and b) sodium building light.

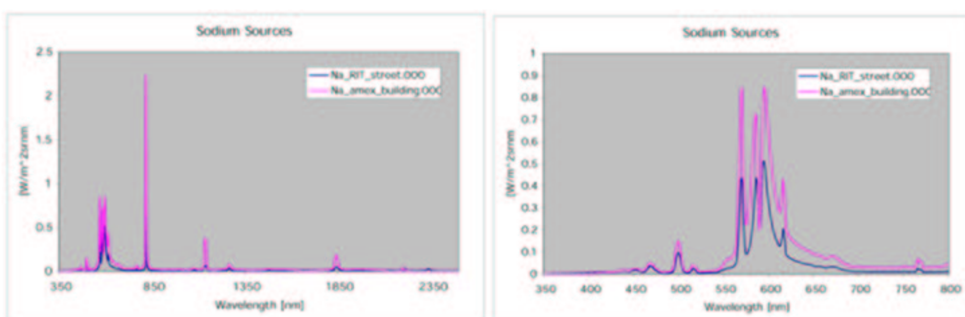

Figure 47: Distribution of sodium sources a) across spectrum and b) in the VIS.

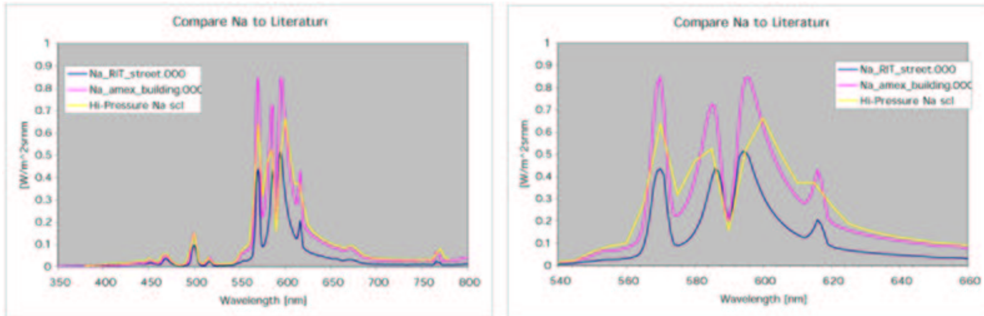

Figure 48: Comparing sodium to published data a) across spectru and b) in the VIS.

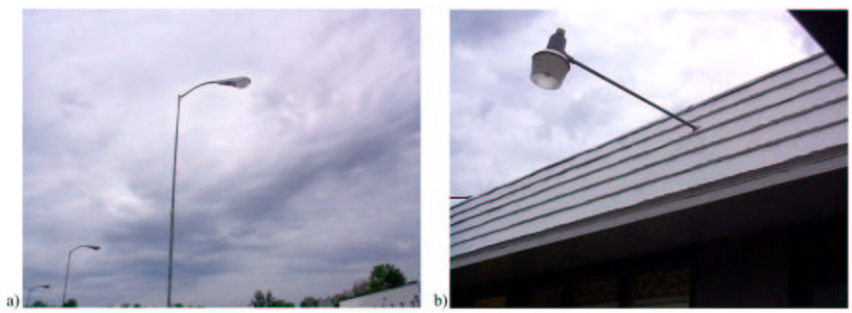

Figure 49: a) Mercury streetlight and b) mercury building light.

The sources were then compared to published data (source: Hunt, R., Measuring Color). The results of this can be seen in Figure 48. Since the published data range from 380-780 nm, only a visible comparison was made. It is evident that the emission peaks correlated very well with literature.

6.1.2 Mercury Lamps

Various types of mercury gas discharge lamps were also measured. These included a standard streetlight and two lamps from the side of a building, illuminating large parking areas. Daylight images of these sources can be seen in Figure 49 and Figure 50. The corresponding spectral distributions can be seen in Figure 51. The spectrum consists mainly of a series of lines, the more prominent of which are at wavelengths of 253.7, 365.4, 404.7, 435.8, 546.1, and 578.0 nm.

Figure 50: Mercury parking lot light.

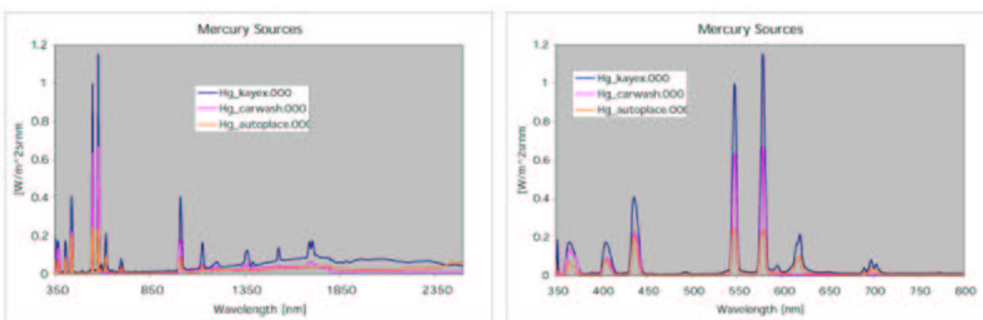

Figure 51: Distribution of mercury sources a) across spectrum and b) in the VIS.

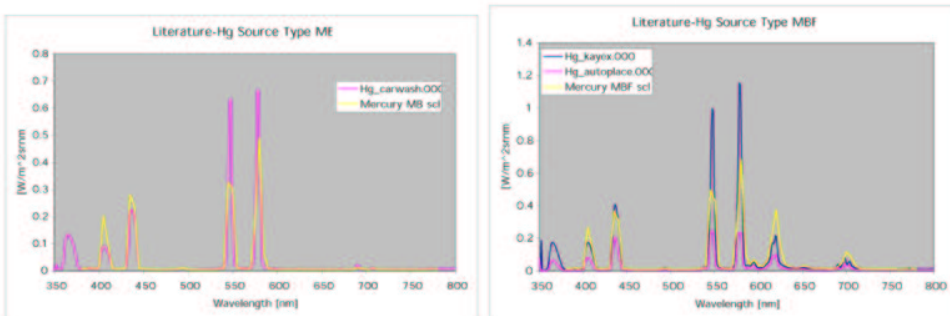

Figure 52: Typical high-pressure mercury lamp type a) MB and type b) MBF.

When comparing these sources to published data, it was found that one of the parking lot sources was a high-pressure mercury lamp type MB (see Figure 52). The other parking lot source was of type MBF and had the same spectral distribution as the streetlight. Both of these can be seen in Figure 52 along with published data. The MBF type has a red-emitting phosphor coated on the inside of the envelope. This improves the color rendering of the lamp appreciably.

6.2 Fluorescent Sources

A typical ceiling fluorescent source was also measured. This measurement was performed with and with out the diffuser over the light. Similarly, another device (PR650) was used to cross-reference the ASD measurements. The first set of ASD measurements for the light source can be seen in Figure 53.

For the most part the distributions are identical. That is, the diffuser has no effect on the output other than to diffuse the light. There is one region, however, located around 365 nm where the diffuser seems to absorb the fluorescent emission line (see Figure 54). This is a very small peak, relative to the entire spectrum, and could be neglected.

The PR650 was also used to measure the spectral output of the source. The results of this can be seen in Figure 55. Here we see that the diffuser simply attenuated the signal. We don't see the 365 nm absorption feature because the range on the instrument is from 380-780 nm.

We then compared the ASD reading to that of the PR650. The results of this can be seen in Figure 56. Both instruments recorded similar spectra.

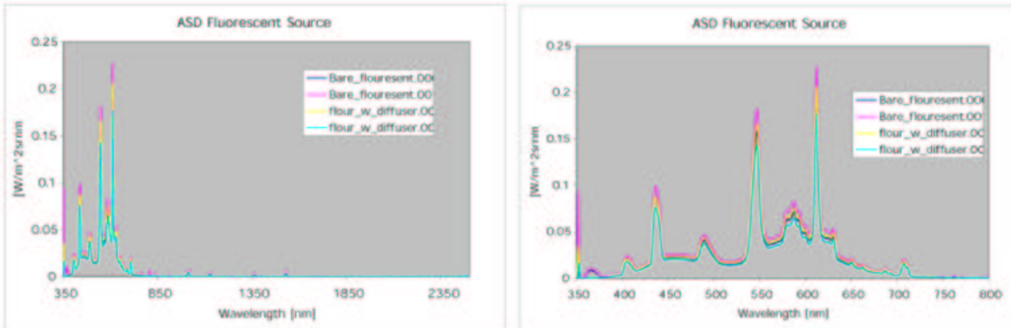

Figure 53: Fluorescent source with and with out diffuser a) across spectrum and in b) VIS region.

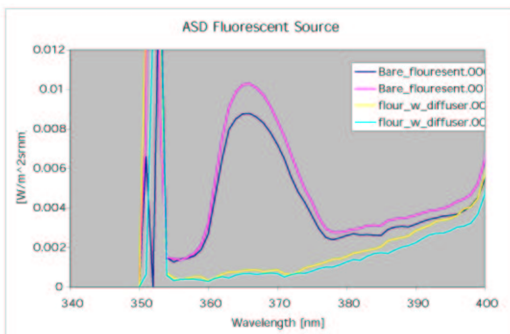

Figure 54: Fluorescent source with diffuser absorbing emission line.

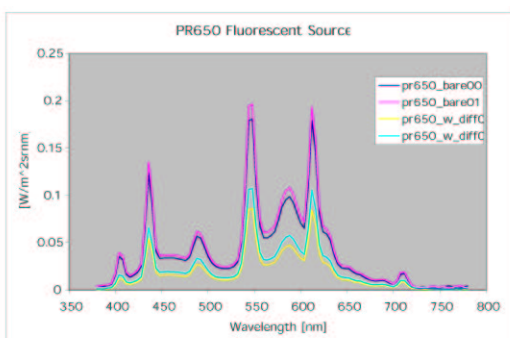

Figure 55: Fluorescent source with and with out diffuser in VIS using PR650.

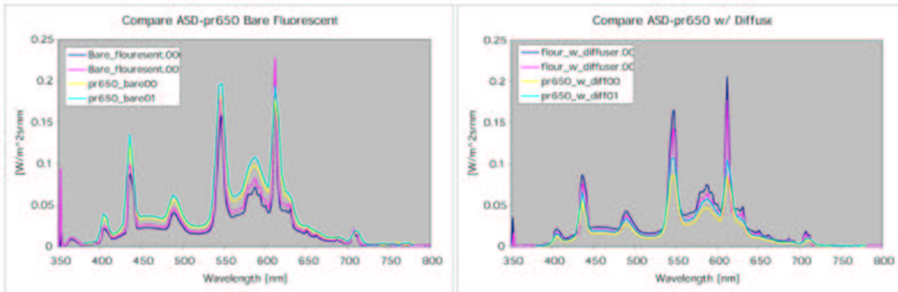

Figure 56: a) Bare and b) diffused fluorescent with ASD and PR650.

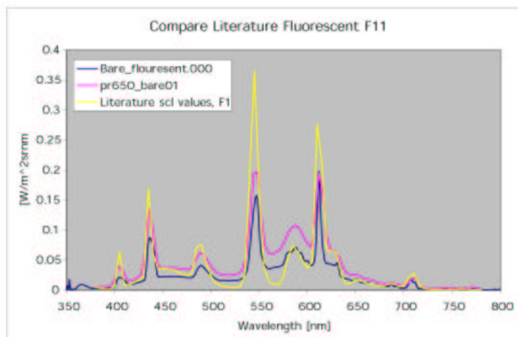

Figure 57: Comparison of published data to measured data.

The fine structure in the ASD readings is evidence of the higher resolution in the instrument (1nm) while the resolution of the PR650 data was 4nm.

Finally, the measured fluorescent data for both the ASD and the PR650 was compared to published data. This comparison can be seen in Figure 57. The spectra looks fairly spikey like the mercury gas-discharge lamps. This is because the lamps consist of a glass tube containing low-pressure mercury gas, in which a gas-discharge is produced. The inside of the tube is coated with phosphors that are excited by the UV lines of the mercury spectrum, particularly that at 253.7 nm, to produce additional light. Therefore the light from these lamps comes partyly from the gas-discharge, but mainly from the phosphors.

There are 3 classes of fluorescent lamps designated *normal*, *broad-band*, and *three-band*. The ceiling lights recorded here are of the three-band type, specifically type F11. This is evident by the high correlation between the

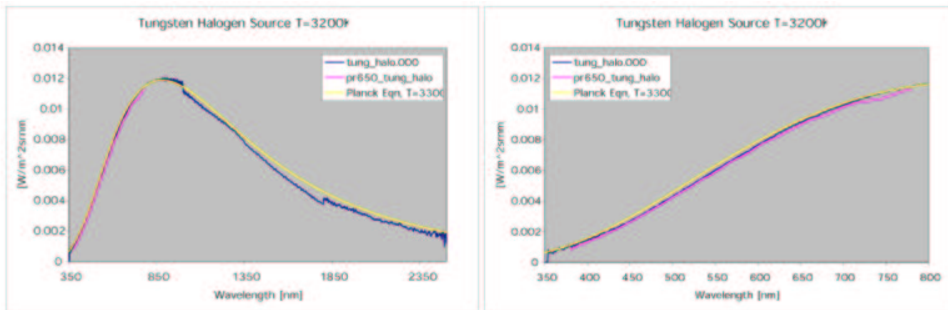

Figure 58: Tungsten-halogen light source a) across spectrum and b) in the VIS region.

published and measured emission lines, as seen in Figure 57. As the name implies, the emission of three-band sources tend to be concentrated in three bands of the spectrum. These bands are also quite narrow and are designed to occur around 435, 545, and 610 nm.

6.3 Incandescent Light Sources

6.3.1 Tungsten-Halogen Lamps

A tungsten-halogen source was measured using both the ASD and the PR650. The results of these measurements can be seen in Figure 58. Incandescent sources are not spiky like some of the earlier gas-discharge or fluorescent sources. This is because in solids and liquids, the atoms are much more closely packed than in gases.

The PR650 almost recorded an exact match to that of the ASD. Additionally, the PR650 recorded a color temperature of 3200 K. For reference, an analytic planckian function was plotted for comparison with a color temperature of 3300 K.

A Appendix A

A.1 Derivation of Irradiance from an Area Source of Radiance, L

The irradiance from and extended area source is simply a modified version of the inverse square law for a point source. The inverse square law is definition as

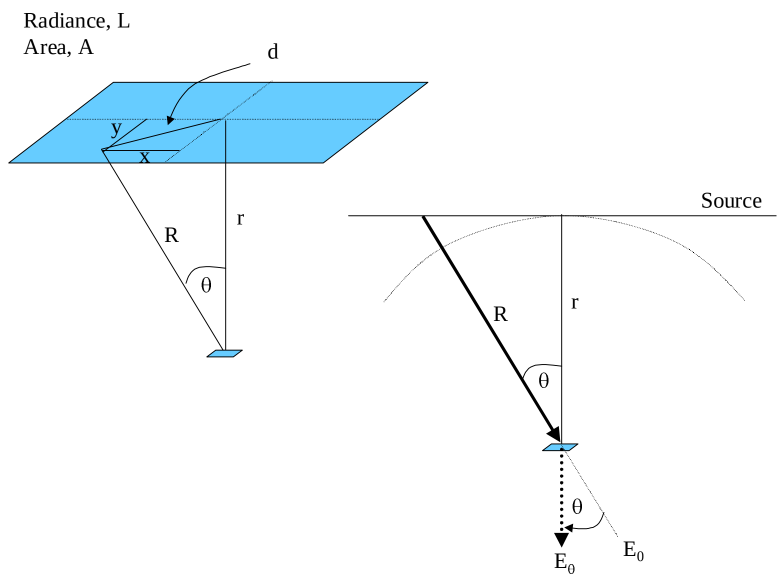

Figure 59: Geometry for calculating the irradiance from an area source.

$$E_2 = \frac{E_1\, r_1^2}{r_2^2} \quad \left[\frac{W}{m^2}\right] \tag{33}$$

which simply states that the irradiance E_2, from a point source varies inversely with the square of the distance, r, from a point source. A similar relationship can be derived for a source of known radiant intensity, I. A few radiometric substitutions of equation (34) yields the relation

$$E = \frac{I}{r^2} \tag{34}$$

which is not a definition but, again, reflects the same variation in irradiance from a point source. We now consider the irradiance from an extended area source. This geometry is illustrated in Figure 59 where we have a source of radiance L, and wish to know the irradiance at some distance. By definition

$$L = \frac{dI}{dA\cos\theta} \quad \left[\frac{W}{m^2 sr}\right] \tag{35}$$

Rewriting equation (34) in differential form and substituting equation (35) into equation (34) yields

$$dE = \frac{L\ dA \cos\theta}{r^2} \tag{36}$$

which, when $\theta = 0$, is the on-axis irradiance for a simple point source.

We also have to consider the off-axis contributions as well. The new path length, R, is longer than, r, as illustrated in Figure 59. Therefore,

$$dE_o = \frac{L\ dA \cos\theta}{R^2} \tag{37}$$

Additionally, we have to take into account projected area effects.

$$dE_\theta = dE_o \cos\theta \tag{38}$$

$$dE_\theta = \frac{L\ dA \cos^2\theta}{R^2} \tag{39}$$

After integrating, factoring in the transmission, and including the source shaping function, we have

$$E = \frac{L\ dA \cos^2\theta}{R^2}\, \tau \cos^m \alpha \qquad \left[\frac{W}{m^2}\right] \tag{40}$$

This is the form of the equation that is used in the DIRSIG ray-tracer.

Additionally, we can rewrite this equation in rectilinear coordinates. We simply need to rearrange the angle term. By geometric inspection of Figure 59 we have

$$\cos^2\theta = \frac{r^2}{R^2} = \frac{r^2}{r^2 + d^2} = \frac{r^2}{r^2 + x^2 + y^2} \tag{41}$$

$$dE_\theta = \frac{L\left(\frac{r^2}{r^2+x^2+y^2}\right)}{r^2 + x^2 + y^2} dA \tag{42}$$

$$dE_\theta = \frac{Lr^2}{\left(r^2 + x^2 + y^2\right)^2}\ dxdy \tag{43}$$

We now must integrate over the source area as a function of x and y. Additionally, this is all performed spectrally. After factoring in the transmission and shaping function we have

$$E = r^2 \cos^m \alpha \int_\lambda \int_x \int_y \frac{L(\lambda)\tau(\lambda)}{\left(r^2 + x^2 + y^2\right)^2}\ dx\ dy\ d\lambda \quad \left[\frac{W}{m^2}\right] \tag{44}$$

References

B References

[1] White, R.A., "Validation of Rochester Institute of Technology's (RIT's) Digital Imaging and Remote Sensing Image Generation (DIRSIG) Model-Reflective Region", MS Thesis, Rochester Institute of Technology, Center for Imaging Science, July (1996).

[2] Kraska, T.A., "DIRSIG: Digital Imaging and Remote Sensing Image Generation Model: Infrared Airborne Validation and Input Parameter Analysis," MS Thesis, Rochester Institute of Technology, Center for Imaging Science, April (1996).

[3] Ientilucci, E.J., "Synthetic Simulation and Modeling of Image Intensified CCD's," MS Thesis, Rochester Institute of Technology, Center for Imaging Science, April (1999).

[4] Weidner, V. R., Hsia, J.J., "Reflection properties of pressed polytetrafluoroethylene powder," Journal of the Optical Society of America, vol. 71, p. 856, July (1981).

[5] Coltman, J. W., J. Opt. Soc. Amer., vol. 44, pp. 468-471, (1954).