7 form•Z RenderZone RadioZity

Contents

7 form•Z RenderZone RadioZity

7.0 Introduction

form•Z RenderZone RadioZity is the version of **form•Z RenderZone** that includes radiosity based rendering, provided by the LightWorks® radiosity engine. With **form•Z RenderZone RadioZity**, the lighting conditions in a scene can be accurately simulated and incorporated into a rendering.

Radiosity techniques were first used in thermal engineering to simulate the transfer of heat energy between surfaces. With the advance of computer graphics in recent years, radiosity has also been used to simulate the transfer of light between surfaces. In fact, radiosity simulates the distribution of light in an environment in a physically accurate manner. The representation of this simulation, called the ***radiosity solution***, describes the distribution of light from all sources across all surfaces in the scene, and it also accounts for light that bounces off surfaces and illuminates other surfaces. It should be noted that radiosity is not by itself a rendering algorithm, such as raytracing or z-buffer. It is an accurate representation of the illumination model of a scene. After such a model has been derived, it is rendered using a raytrace or z-buffer rendering procedure. In **form•Z RenderZone RadioZity** this is done by applying **Shaded Render*** or one of the **RenderZone*** rendering modes to a radiosity solution. A radiosity solution can also be used to produce an interactive shaded rendering using **QuickDraw 3D*** on the Macintosh and **OpenGL*** on Windows.

Non-radiosity based rendering methods use a simple illumination model to simulate the interaction of light with the environment. During such a rendering, the light intensity at a pixel is calculated by shooting a ray from the pixel's position back to a light source. Taking into account the falloff of the light intensity, the distance of the pixel from the light source determines the light intensity at the pixel. If this ray is blocked by a surface which is located between the pixel and the source, the pixel is in shadow and no illumination of the pixel occurs. A scene rendered with this approach only illuminates pixels which are directly hit by a light source. In order to give all other pixels some light intensity, a constant global ambient light is added. This ambient light is a crude approximation of the light reflecting off surfaces, illuminating other surfaces indirectly. Areas in a scene which are illuminated by the constant ambient light alone show no variance in intensity, which is responsible for the artificial appearance of such renderings.

The radiosity solution correctly simulates this ambient light by computing light distribution which results from light reflecting off surfaces. It is this secondary illumination that gives radiosity based renderings their high degree of realism.

The radiosity solution is ***view independent***. This means that, once a radiosity solution has been completed, any number of renderings can be produced from it without having to calculate the light intensities again. As a result, radiosity based renderings are significantly faster than renderings using the simple illumination model, given the same amount of geometry to process.

7.0.1 What is radiosity and how does it work

Formally, ***radiosity*** is a measure of energy leaving a surface, per unit area, per unit time. The underlying strategy of radiosity is to divide all surfaces in a scene into small polygons, called ***patches*** and to compute the radiosity value for each. The radiosity value of all patches is determined by first distributing energy from all primary light sources to all patches. Each patch absorbs some of the energy it receives and reflects the rest. The amount of reflection depends on the diffuse reflection parameter associated with the patch. After the distribution of light from the ***primary sources***, the patch which reflects the most amount of energy is now considered a light source itself, called a ***secondary source***. It distributes its reflected energy to all other patches in the scene. Next, the patch which reflects the second highest amount of energy is processed, and so on. Note that the first patch may reappear at some point in the list of light emitting patches, because it has received enough energy from other patches to become the brightest patch in the scene. Obviously, this process will continue forever, since surfaces typically emit a portion of the energy they receive. It is therefore necessary to specify some criteria, which will stop the radiosity process. These criteria may be a given amount of time that the process is allowed to run, a given number of cycles, or the amount of energy absorbed by a patch relative to the overall energy emitted by all primary sources.

By the time the radiosity process stops, a mesh of small patches with varying intensities has been generated. This mesh can be rendered in the standard way, substituting the per pixel intensity calculation of the simple illumination model with the precalculated intensity of the patch. The rendering is also able to smooth the intensity changes between neighboring patches, giving the impression of a gradual change of illumination and soft shadow boundaries. Such a rendering will show a varying amount of light in areas which are not directly illuminated by the primary light sources. In contrast, a rendering generated with the simple illumination model shows no such variation. A scene rendered with a simple illumination model and from a radiosity solution is shown in Figure 7.0.1.1.

a

b

Figure 7.0.1.1: A scene rendered with (a) a simple illumination model and (b) from a radiosity solution.

7.0.2 Reflections in radiosity

Photorealistic renderings typically simulate three kinds of reflections: ***diffuse*** and ***specular*** reflections of light, and ***mirrored*** reflections of the environment. Radiosity simulates diffuse reflections. Specular and mirrored reflections may be applied when using raytracing to render the radiosity solution.

Diffuse reflection of light is responsible for the gradual change of shading on a curved surface. A light ray which hits a surface at a given angle is scattered by the surface and an infinite number of new rays is emitted at equal intensity in all directions, as shown in Figure 7.0.2.1 (a). The intensity of the reflected rays is based on the angle of the incoming ray and on the ability of the material to produce diffuse reflections. In addition, the color of the reflected rays is composed of the incoming color filtered through the color of the surface material, generating what is commonly perceived as the color of a material. Surfaces which show effects from diffuse reflection alone will result in a dull, matte appearance. Figure 7.0.2.1 (b) shows a spherical surface which exhibits only diffuse reflection.

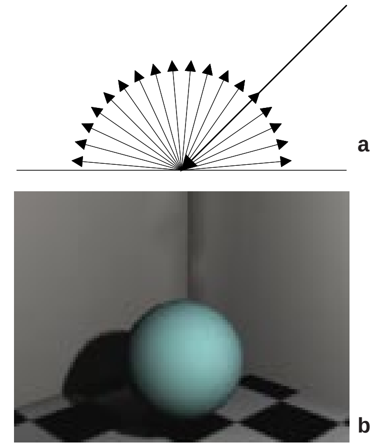

Figure 7.0.2.1:
A diffuse reflection (a) diagram and (b) rendering.

Specular reflection of light bounces the incoming light ray off the surface at the same, or close to the same angle, as illustrated in Figure 7.0.2.2 (a). That is, specular reflection can only be observed, if the outgoing ray exactly hits the eye. The reflected light ray of a specular reflection does not take on the color of the material, but maintains the color of the incoming light. As a result, specular reflections generate hotspots on curved surface in the color of the light, which is typical for highly polished and glossy materials. Figure 7.0.2.2 (b) shows a spherical surface which has a high degree of specular light reflection.

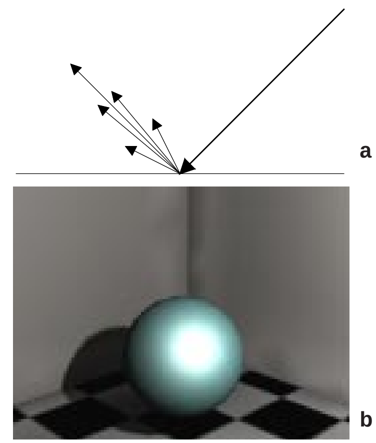

Figure 7.0.2.2:
A Specular reflection (a) diagram and (b) rendering.

Mirrored reflections of the environment are generated by reflecting the line of sight off a surface at the incoming angle and intersecting the reflected ray with whatever surface it hits next. Mirrored reflections of the environment are commonly simulated in computer graphics through raytracing techniques. Figure 7.0.2.3 shows a sphere with mirrored environment reflections.

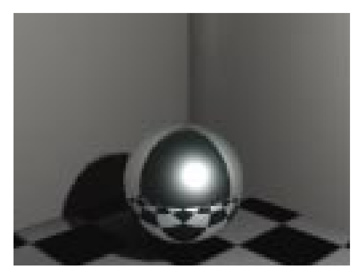

Figure 7.0.2.3: Mirrored environment reflections.

The final color of a pixel on a surface that exhibits reflections is a mixture of the colors generated by all three types. As already noted, a radiosity based rendering receives its diffuse contribution from the precalculated intensities at the mesh patches, whereas specular and mirrored reflections, transparencies, and other rendering effects are calculated by the same algorithms used by the simple illumination model.

7.0.3 Performance of radiosity solutions

Computing a highly accurate radiosity solution can be an intensive task. Much effort is spent calculating the intensities at the radiosity mesh patches. On a given patch, a number of points are chosen randomly. From each point a light ray is cast to the vertices of all other patches in the scene. The illumination of all patches by a single source patch is called an ***iteration***. The distance from the sampled point to the vertices determines the amount of energy received at the other patches. Recall that light intensity falls off with the square of the distance the light ray travels. In order for every patch to illuminate all other patches in the scene, ***n*** * ***n*** iterations must be performed, where ***n*** is the number of patches in the scene. This is an exponential process. In other words, doubling the number of patches in the scene results in four times the number of iterations. It is apparent that a dense mesh of several thousand patches will significantly increase the time required to process patch illumination. Therefore, minimizing the amount of patches is a primary concern.

The accuracy of a radiosity solution is measured by the percentage of the light that has been absorbed by the radiosity patches over the total light initially emitted by all the active light sources. Since the computation of the radiosity solution is not a linear process, a solution, which is for example 70% accurate, can be achieved in a relatively short period of time, whereas the next 10% will take considerably longer. As the solution approaches 100% accuracy, increasingly more time will be required. In many cases a solution that is about 80% complete will generate renderings of acceptable quality. It is only for solutions which require a very high degree of accuracy that a great amount of time will be required.

The density of the solution mesh directly affects the quality of a rendering. Solutions with a low mesh density generate accurate results faster and can be rendered faster, but will most likely exhibit visible artifacts at areas of sharp intensity contrast. Low density radiosity solutions are appropriate for real time renderings such as those produced by **OpenGL** on Windows or **QuickDraw 3D** on the Macintosh. Solutions that are based on a high mesh density will take longer to process but will reduce or eliminate visible artifacts. Such solutions are typically used for photorealistic, raytrace, or z-buffer renderings. Both scenarios are offered by the radiosity implementation of **form•Z**.

7.1 Rendering with radiosity

form•Z RenderZone RadioZity offers four menu items in the **Display** menu (Figure 7.1.0.1) for producing radiosity solutions and renderings. These are supported by dialogs where the different options can be set. In addition, two types of lights that are unique to radiosity are available and all lights are specified with more accurate parameters.

Radiosity is a mode within **form•Z**. Once you enter the radiosity mode, you cannot make changes to the model until you exit. While in the radiosity mode, you can adjust the radiosity parameters, change views, produce radiosity illuminated renderings (using **QuickDraw 3D*** on the Macintosh, **OpenGL*** on Windows, **Shaded Render***, or **RenderZone***) and save the **form•Z** project.

A radiosity session is typically executed in an iterative manner, where the solution and rendering are gradually refined until the desirable quality is reached. Following is an outline of a typical radiosity session:

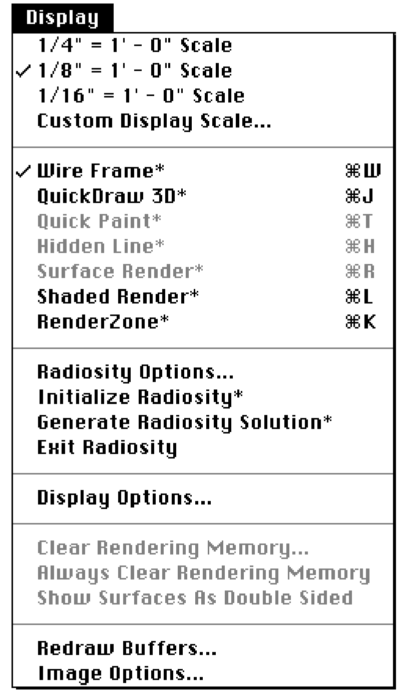

Figure 7.1.0.1: The **Display** menu.

- An interior scene is modeled, using multiple light sources.
- The **Radiosity Options**... menu command is selected.
- This invokes the **Radiosity Options** dialog, where the desired parameters are chosen. Initially, the density of the mesh may be set to a coarse resolution, so that the radiosity solution may progress quickly and it is not necessary to wait a long time to check for proper lighting conditions in the scene.
- After clicking **OK** in the **Radiosity Options** dialog, the **Initialize Radiosity*** item is selected and **form•Z** enters the radiosity mode.
- Next the **Generate Radiosity Solution*** menu item is selected.
- A progress bar appears, indicating the density of the mesh and the progress of the solution towards a desired level of completion.
- At specified intervals a shaded preview is drawn, allowing you to visualize the solution.

If it turns out that some light sources are not at the desired intensity, the radiosity progress is canceled by selecting **Exit Radiosity**. Changes to the intensity of the light sources are made and **Generate Radiosity Solution*** is executed again. As the preview renderings of the solution show increasingly satisfactory results, it is possible to set the mesh density parameters to a higher resolution, so the image quality may improve as well. This process of refinement may be repeated several times until the preview images show a satisfactory result. After the solution has reached the desired level, several renderings from different views may be produced. The images can be saved using the **Save As...** command in the **File** menu.

- The radiosity work is completed by selecting the **Exit Radiosity** menu item.

While in radiosity mode, it is possible to make changes to both surface styles and lights. After editing a surface style, the radiosity solution may be continued, and from that point on it will use the new parameters of the surface style. For example, if the color of a surface style was changed, the light reflected off faces using this surface style will now be filtered through the new color. Note, that the light which was distributed in the scene prior to the editing of the surface style is not affected by the color change. Editing a surface style during a radiosity process is intended to make subtle changes to its parameters to fine tune the appearance of surfaces. In general, small changes to a surface style will not affect the accuracy of the solution. At the same time, there is nothing to prevent the user from drastically altering a surface style, which would result in a situation where the appearance of a surface and the corresponding lighting conditions of a scene do not correspond to the real conditions.

In a similar manner, lights can also be edited while in radiosity mode. Unlike surface styles, changes made to a light source requires that the current radiosity solution be deleted before the radiosity process is continued. Any changes made to the lights in the scene will cause the radiosity process to start over as soon as **Generate Radiosity Solution*** is selected from the **Display** menu. The ability to edit lights while in radiosity mode is intended to be a shortcut. Instead of requiring you to select the **Exit Radiosity** menu command before editing a light, the lights can be changed directly. To prevent accidental changes to a light while in radiosity mode, a warning message is posted before the changes are accepted, which gives the user the option to cancel the changes.

7.1.1 The radiosity commands

Radiosity Options...

Selecting this item invokes the **Radiosity Options** dialog discussed in section 7.1.3 and shown in Figure 7.1.3.1. This dialog, which can also be invoked from the next two menu items, contains the radiosity parameters and settings.

Initialize Radiosity*

When this menu item is selected, **form•Z** enters the radiosity mode. Initializing radiosity constructs mesh patches from all of the objects in the scene based on the parameters in the **Radiosity Options** dialog. The **form•Z** project is frozen at this point and no changes can be made to the model. In addition, the content of the scene cannot be changed and objects and layers cannot have their visibility changed. To indicate that the radiosity mode is active, the **form•Z** environment dims all unavailable tools, menu items, and palettes as shown in Figure 7.1.1.1.

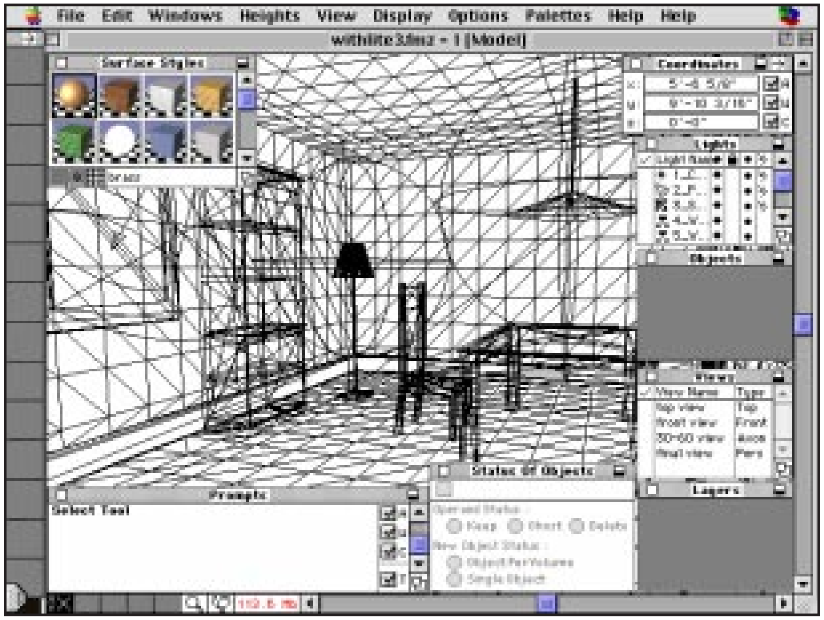

Figure 7.1.1.1: The **form•Z** environment when radiosity is initialized.

Once the **Initialize Radiosity*** command has been selected, it can be chosen again at any time to, for example, correct some of the parameters set in the **Radiosity Options** dialog. When selected a second time before the solution is started, the initial mesh is simply regenerated. After the solution has been started, the current solution is discarded first before the solution is initialized again.

Since a radiosity solution is view independent, all of the tools and menu items relating to view control remain available. The lighting model that is represented in the radiosity solution can not be used in **Quick Paint***, **Hidden Line***, and **Surface Render***. These items appear dimmed in radiosity mode.

Generate Radiosity Solution*

When this item is selected, the radiosity process is started. If **Initialize Radiosity*** was executed prior to selecting **Generate Radiosity Solution***, the radiosity process is started immediately. Otherwise radiosity is initialized first. Starting the radiosity solution invokes continuous iterations of patch illumination. A progress bar appears indicating how many mesh polygons currently exist in the scene and how far the solution has progressed toward the user defined goal. At a given interval, a shaded rendering or a wireframe rendering show the current state of the solution. The shaded rendering uses the calculated patch intensities as the base for the diffuse reflection for the objects in the scene. The execution of the radiosity solution can be stopped at any time by clicking on the **Stop** button of the progress bar, or by hitting the shortcut key combination assigned to the **Stop** action.

Exit Radiosity

This menu item is initially dimmed. It becomes active and can be selected after the radiosity mesh has been initialized by choosing the **Initialize Radiosity*** or **Generate Radiosity Solution*** commands. When selected, the program exits the radiosity mode and all tools, palettes, and menu items become active again.

After a radiosity solution has been stopped by pressing the **Stop** button or by hitting the cancel shortcut key, it is possible to restart the solution by selecting the **Generate Radiosity Solution*** menu item. Typically, a solution is stopped in order to change some of the parameters in the **Radiosity Options** dialog or to fine tune the solution based on the visual feedback provided by the preview renderings and by the information shown in the progress bar. The solution may also be stopped when it is considered sufficiently advanced, in order to generate final renderings or real time renderings. The display commands, **Wireframe***, **Shaded Render***, **RenderZone*** and **QuickDraw 3D*** or **OpenGL***, remain active during radiosity. They can be selected to produce a rendering using the radiosity solution. It is therefore possible, for example, to set up a new view in **Wireframe*** display and render a **RenderZone*** image, or select **QuickDraw 3D*** or **OpenGL*** and navigate through the scene in real time using standard view manipulation tools.

When a radiosity solution is interrupted, it is possible to preserve it at the current state, so that it may be continued later. This is done by writing it into a temporary file when **Exit Radiosity** is executed or when toggling between projects. The solution can also be saved more permanently by saving it with the project when **Save** or **Save As...** is executed. Whether or not these options are active is set in the **Radiosity Preferences** dialog (Figure 7.1.2.1), which is invoked from the **Preferences** dialog.

7.1.2 The radiosity preferences

The **Preferences** dialog of **form•Z RenderZone RadioZity** contains a button for selecting certain radiosity options.

Radiosity...: Clicking on this button invokes the **Radiosity Preferences** dialog (Figure 7.1.2.1).

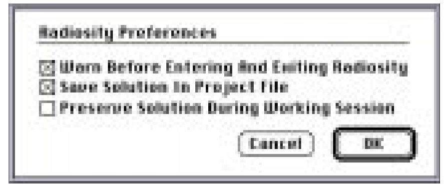

Figure 7.1.2.1:
The **Radiosity Preferences** dialog.

Warn Before Entering And Exiting Radiosity: When this option is selected, a warning message is posted when one of the following conditions exist:

- The **Generate Shadows** option is on in the **Radiosity Options** dialog, but none of the active lights have their shadows attribute on.
- A distant light is active, which does not cast shadows.
- Radiosity is exited by selecting **Exit Radiosity** from the **Display** menu, a window is closed, or you switch to another project, while **Preserve Solution During Working Session** is off (see below).

Save Solution In Project File: When this option is on and **Save** or **Save As...** is selected from the **File** menu the radiosity solution is saved with the **form•Z** project. When such a file is opened, and **Generate Radiosity Solution*** is selected from the **Display** menu, the previously saved radiosity solution is restored and the radiosity process continues from where it was interrupted when the file was saved. This option is on by default.

Preserve Solution During Working Session: When this option is selected, the current radiosity solution is not discarded when selecting the **Exit Radiosity** command from the **Display** menu, but is saved to a file on disk. If no change is made to an object or light and **Generate Radiosity Solution*** is selected from the **Display** menu, the radiosity process continues from where it was interrupted. However note that, if the geometry of an object or a light's parameters change, the radiosity solution also has to change and can not be continued, regardless of whether this option is on or off. This option also affects whether a radiosity solution will be preserved to be continued later when we switch active projects. That is, given that only one radiosity solution may be in progress at any time, which is the solution of the active project, when we switch active projects, the radiosity solution of the previous project (when there is one in progress) cannot be continued. It either has to be discarded or saved on disk. With this option on, the solution is saved to disk and is restored when the project becomes active again. If this option is off, the current radiosity solution is always discarded when selecting **Exit Radiosity** or when switching projects. Saving the solution on disk may require some noticeable time. Consequently, this option is off by default.

Note that, when **Preserve Solution During Working Session** is not selected, it is not possible to switch between active projects. A warning message will be posted after selecting **Exit Radiosity**, which allows the user to avoid accidental deletion of the radiosity solution.

7.1.3 The radiosity parameters

This section discusses in greater detail the parameters which influence the visual quality of a radiosity solution. The radiosity process lends itself to a continuous refinement. Initially a coarse mesh may be specified to arrive at an accurate solution quickly. As the scene is fine tuned, a denser mesh yields better images, at the cost of longer processing time. The values, which determine the mesh density and other radiosity related parameters, are set in the **Radiosity Options** dialog, shown in Figure 7.1.3.1.

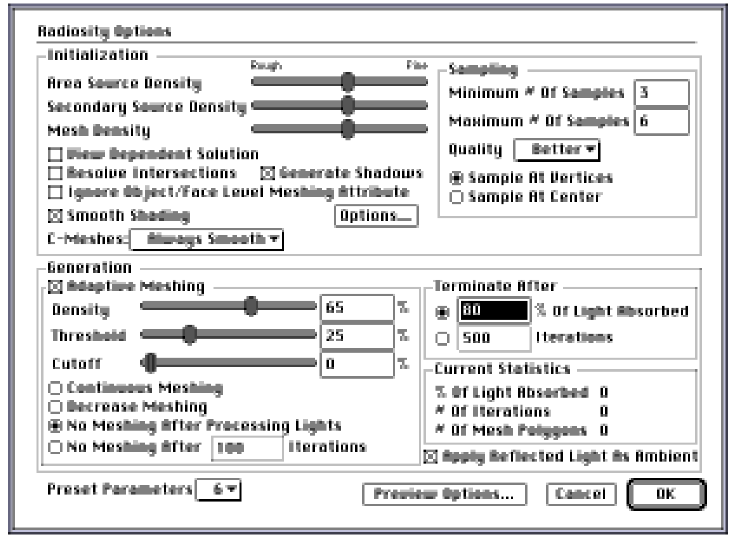

Figure 7.1.3.1: The **Radiosity Options** dialog.

This dialog is divided into two parts, **Initialization** and **Generation**. The parameters in the **Initialization** category are used by the **Initialize Radiosity*** menu command, whereas the parameters in the **Generation** category affect the radiosity process executed by the **Generate Radiosity Solution*** menu command. Recall that the **Generate Radiosity Solution*** command can be used to both initialize and execute the radiosity solution, if selected without initializing the solution first.

Initialization

The mesh, which is the base for a radiosity solution, consists of three separate parts whose densities may be set independently from each other. The first part describes the geometry of all area light sources, called ***primary patches***. The second part describes all the objects in the scene which are not active light sources, called ***secondary patches***. The patches representing the area lights and secondary sources are generated when the radiosity process is initialized and their number stays constant from then on. The third part is the actual ***radiosity mesh***, which is a finer subdivision of the primary and secondary patches and may grow in density as the solution progresses. The light intensities that are calculated by the radiosity process are recorded at the points of this radiosity mesh, which is also referred to as the ***recording mesh***.

Area Source Density: This slider bar determines the density of the primary patches. The minimum number of patches for area lights is generated, when the slider bar is moved all the way to the left. The highest number of area light patches is created by moving the slider bar all the way to the right.

Secondary Source Density: This slider bar is used to specify the density of the secondary patches. The minimum number of patches for surfaces which do not actively emit light is generated, when the slider bar is moved all the way to the left. The highest number of secondary patches is created by moving the slider bar all the way to the right.

Mesh Density: This slider bar can be used to specify the density of the radiosity mesh at initialization. When it is at the left end, the mesh density is equal to the number of primary and secondary patches. As the slider is moved to the right, the mesh density increases.

While the primary and the secondary patches are able to emit light, the subdivided mesh is only able to record light intensities. In general, it is not necessary to create a high density for the patches which emit light. However, to show sharp shadow boundaries or a rapid falloff in light intensity on a surface, the mesh which records the light intensities may have to be very dense to show visually acceptable results. The additional density of the recording mesh can be achieved in two ways. First, the **Mesh Density** slider bar can be used to create an initially high mesh density. Second, the **Adaptive Meshing** option may be used to increase the mesh density during the execution of the radiosity solution in areas of strong intensity contrast. During a single iteration of the radiosity solution, light is emitted from several points of a primary or secondary source patch and is recorded at each mesh polygon. There is a direct relationship between the density of primary and secondary patches and the quality of the illumination in a scene. The denser the source patches are, the higher the quality of a rendering will be and the more computing time is required.

Resolve Intersections: When this option is selected, patches which intersect each other are split into separate patches. When the object that casts shadows intersects the surface on which the shadow is cast, it is possible that the shadow or light may leak out from underneath the former object. Such artifacts are shown in Figure 7.1.3.2 (a). At the front right leg, one can observe that light leaks from underneath the leg into the area where there should be shadow. When this option is on, the intersecting surfaces are split, which causes the mesh to align with the edges of the object and eliminates the light leaks (Figure 7.1.3.2 (b)).

Note that when faces only touch each other, they are not split. Because this may still result in leaks, it is sometimes advisable to make a few adjustments, such as to insert a segment where the face touches or to slightly move the objects so that they intersect.

a

b

Figure 7.1.3.2: Rendering with **Resolve Intersections** (a) off and (b) on.

View Dependent Solution: If only a single rendering is required from a radiosity solution, it is not necessary to generate dense areas of the radiosity mesh in regions that are not visible. Turning this option on will avoid additional mesh subdivision in invisible regions of the scene, which leads to shorter execution times of the solution. While it is still possible to render images from different views, optimal results are only achieved from the view which is set when the radiosity solution is executed.

Generate Shadows: When this option is selected, the radiosity solution incorporates shadows cast from light sources, which have their respective **Shadows** option turned on. Recall that the **Shadows** option for individual light sources is set in the **Light Parameters** dialog. Shadows are only cast by objects which have their shadow casting attribute on and are cast only on objects which have their shadow receiving attribute turned on. Both attributes can be set using the Set Rendering Attributes or Query tool.

Ignore Face/Object Level Meshing Attributes: Selecting this option disables the meshing attributes applied to an object or face using the Set Rendering Attributes tool. When this option is off, if an object or face has meshing attributes, these override the meshing applied at initialization and generation.

Smoothing...: Clicking on this button invokes the **Smooth Shading Options: Radiosity** dialog, shown in Figure 7.1.3.3. The options in this dialog are the same as for **QuickDraw 3D***, **OpenGL***, **Shaded Render***, and **RenderZone***. They determine which surfaces are smooth shaded when rendered. Note that these rendering types ignore their own smoothing settings, when they are applied to a radiosity solution.

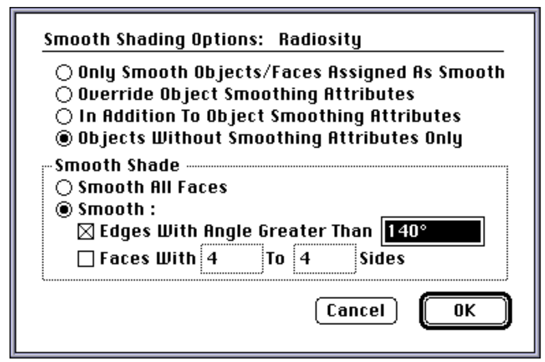

Figure 7.1.3.3: The **Smoothing Options** dialog.

Sampling parameters

The **Sampling** group of parameters determine the quality of the illumination performed during radiosity processing.

Minimum # Of Samples, Maximum # Of Samples: When light leaves a primary or secondary source patch, it is emitted from a number of randomly chosen points on the patch surface. The minimum and maximum number of sample points are entered in these fields. Note that large values entered in these fields may unnecessarily extent the progress of a solution. As a general guideline, minimum values between 1 and 5 and maximum values between 10 and 20 create acceptable results at a reasonable speed. How many points are sampled across the patch surface ultimately depends on the item selected from the **Quality** menu (see below) as well as other factors such as the size of a patch.

Quality: This pop up menu contains items which set the radiosity parameters to a predefined level of quality. The quality of a radiosity solution and the speed at which it is derived are inverse proportional to each other. That is, the lowest quality setting derives a solution quickly, but renderings based on such a solution are usually of poor quality. The highest quality setting takes longer to generate a solution, but the quality of images is higher.

Simple: When this item is selected, the minimum number of samples is used.

Better: When this item is selected a number of samples between minimum and maximum is used, but is more often closer to the minimum than the maximum number.

Best: When this item is selected the number of samples used is more often closer to the maximum value.

While the previous options in the **Sampling** group determine the density of rays that leave the light emitting patches, the next two options control how frequently the light intensities are recorded at the receiving patches.

Sample At Vertices: When this option is selected, the light emitted from a primary or secondary source is recorded individually at each of the vertices of a mesh polygon. Given that all mesh polygons are triangles, the light calculation is performed three times.

Sample At Center: When this option is selected, the light emitted from a primary or secondary source is recorded once, at the center of a mesh polygon. Center sampling, while faster, may require a denser mesh to avoid certain illumination artifacts. On the other hand this option generates softer shadows, which decreases the need for a dense mesh at shadow boundaries. Which of the two methods is more suitable depends greatly on the scene. In general, vertex sampling is more accurate, but may generate a denser mesh, whereas center sampling is faster at the expense of some illumination artifacts.

Generation

The **Generation** parameters influence the density of the mesh and the appearance of the solution.

Adaptive Meshing: When this option is turned on, the mesh may be subdivided into a finer resolution than its initial density. To what extent this subdivision occurs is determined by the next three parameters.

Density: This slider bar indicates the maximum density the subdivisions can reach. The further to the right the slider bar is moved, the finer the subdivisions may become. Values of 50% and below will generate coarse subdivisions. Values above 80% should be used with caution, since they may quickly generate an excessive amount of mesh polygons, which may unnecessarily slow down the radiosity solution. Values around 60% are a good compromise between speed and quality, while values between 70 and 80% are appropriate for high quality images. Two pairs of a wireframe and shaded image are shown in Figure 7.1.3.4. The first pair uses a density setting of 60%, whereas the second uses a density of 80%. Note the difference in mesh density and the resulting image quality.

a

b

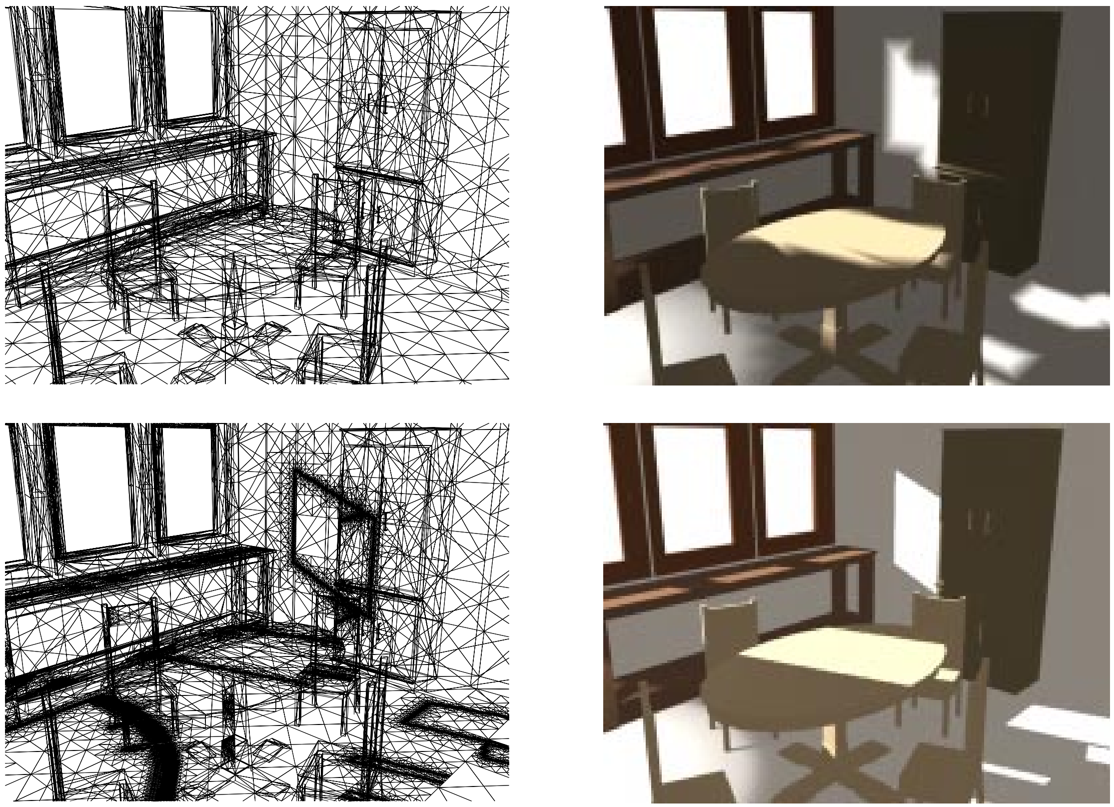

(1) (2)

Figure 7.1.3.4: (1) Wireframes and (2) shaded renderings of a scene using **Density** set at (a) 60% and (b) 80%.

Threshold: This parameter determines when a mesh polygon will be subdivided, due to the difference of intensities at adjacent polygons. If the difference between the intensities of a mesh polygon and all its neighboring polygons is higher than this value, the polygon is subdivided. In other words, the lower the **Threshold** value, the more subdivision will occur.

Cutoff: This parameter can be used to avoid subdivision of mesh polygons in dark areas. Since an image does not show a great amount of illumination detail in underexposed areas, a dense mesh may not be necessary. A low **Cutoff** parameter causes mesh subdivision in dark areas and hence causes a denser mesh. Mesh density may be decreased by increasing the **Cutoff** parameter, which causes areas of increasing brightness to not be meshed.

Adaptive meshing can be specified globally for all objects in the scene or on a per object or face basis. In many cases a dense mesh may be required in a localized area of the scene, while the rest of the scene can achieve satisfactory results at a much lower resolution. This is, for example, the case when a cone light shines on the floor of a room. The mesh polygons around the edges of the cone illumination need to be meshed densely to account for the rapid falloff of intensity in this region. The rest of the scene may be meshed at a much lower resolution than the floor, since the light reflected off the floor only generates very gradual light distribution on the walls and ceiling. Adaptive meshing alone already addresses this issue by meshing the area of sharp contrast more than other areas. The meshing can be even further reduced by assigning to objects or faces, which are known not to require extensive meshing, their own lower meshing parameters. Or vice versa. That is, it is possible to specify low meshing parameters for the entire scene, and then assign higher meshing parameters to objects or faces, which are known to require dense meshing.

Meshing conditions

One of four mutually exclusive options may be used to automatically adjust the sensitivity of the subdivisions as the radiosity solution progresses. A dense mesh is usually only necessary when light from primary sources is absorbed by the mesh, because the light intensity is relatively high and may generate sharp shadow boundaries. Light reflected from secondary sources is typically much lower in intensity and is largely responsible for a soft and gradual falloff of light on other surfaces. A lower density mesh is sufficient to display those intensities. One way to address this issue would be to let a solution progress with a certain set of parameters generating a dense mesh, while primary sources are handled. As soon as this stage is over, one could interrupt the solution, set the meshing parameters to values that will generate a much coarser mesh and restart the solution. This effect can be controlled automatically by selecting one of the following options.

Continuous Meshing: When this option is selected, the mesh parameters stay the same during the entire length of the radiosity execution. While this guarantees the highest quality results, the mesh density may be higher than necessary in areas of low contrast, and the radiosity process may be unnecessarily extended.

Decrease Meshing: When this option is selected, the mesh parameters are proportionally relaxed as the solution progresses. That is, the closer the solution has advanced toward full accuracy, the lower the meshing parameters become. This can be done safely, because towards the end of a solution process, light intensities reflected from secondary sources become increasingly lower, allowing for a lower density mesh. This option is a good compromise between speed of a radiosity solution and quality of the resulting images.

No Meshing After Processing Lights: The decrease of meshing parameters can be taken even one step further. When this option is selected, meshing is turned off completely, after all primary sources have emitted their light. In many cases, this will generate acceptable image quality, at the lowest possible mesh density.

No Meshing After *n* Iterations: Meshing can also be stopped after a given number of iterations. Recall that one iteration is defined by a single source patch, primary or secondary, emitting light to all polygons in the mesh. The value for the number of iterations to be completed before meshing is turned off is entered in the text field.

Terminating conditions

One of two mutually exclusive terminating criteria may be specified. They cause the automatic termination of the radiosity solution when they are reached. In other words, they offer additional control over unnecessarily lengthy processing.

% Of Light Absorbed: When this option is selected, the radiosity calculations terminate after a certain amount of light is absorbed. This amount is expressed as a percentage entered in its field. It indicates how much of the light, which was initially emitted by all primary light sources, must be absorbed by the radiosity mesh before the solution is stopped. Typically, a solution very quickly advances to a level of about 60 - 70% and takes exponentially longer as it progresses towards 100%. It is during the last 20% when the subtle change of illumination of indirectly lit surfaces is generated, which gives radiosity based rendering their realistic look. On the other hand, a solution which is 95% complete may sometimes hardly show any difference from a solution which is 85% complete. While the solution which is 95% complete may take a long time to finish, an 85% complete solution may be generated much faster. It is up to the user to decide how much time can be invested toward arriving at a visually acceptable result

Iterations: When this option is selected, the radiosity calculations will terminate after a certain number of cycles, which is determined by the number entered in its field. While this method does not give an indication of completion, it may be more suitable for determining how long one is willing to let a solution execute. For example, if a solution executes 10 minutes, and during this time 1000 iterations are completed, one can interpolate the time that will be required for a given amount of iterations.

How shadows are generated in radiosity

Recall that, while diffuse reflections are included in a radiosity solution, specular reflections are generated at the time the radiosity solution is rendered. That is, the diffuse illumination of a surface is computed by the radiosity process and is stored with the radiosity mesh, whereas specular illumination is calculated for every pixel when the radiosity solution is rendered. This also applies to shadows. Shadows resulting from the diffuse distribution of light are generated during the calculation of the radiosity solution, while shadows from specular light are generated when the radiosity solution is rendered. Because of this, different options turn on/off the two types of shadows.

Whether the radiosity solution will compute shadows, including diffuse shadows, is determined by the **Generate Shadows** option in the **Radiosity Options** dialog. Whether specular shadows will be generated is determined by the **Shadows** option in the **RenderZone Options** dialog. The latter can be **Hard (Raytraced)** or **Soft (Mapped)** shadows.

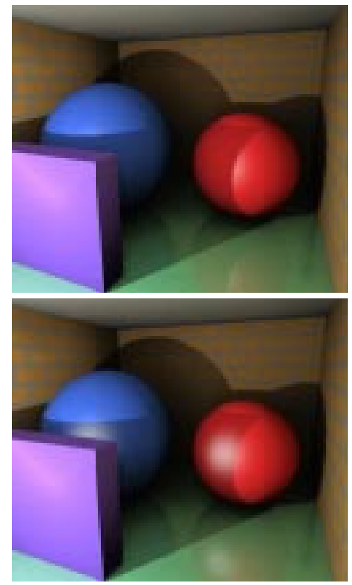

a

b

Figure 7.1.3.5: Specular shadows with **Shadows** option (a) on and (b) off.

In the real world, there is of course no distinction between shadows from diffuse and shadows from specular illumination. Shadows simply occur where there is no light. In a radiosity based rendering the distinction is made because processing of illumination is divided into diffuse and specular computations. As a result, shadows must be calculated for both processes. However, the shadows generated by specular illumination occur only where glossy hotspots are visible on curved surfaces with a high degree of specular reflection. Because images do not generally exhibit a lot of glossy areas, it may be valid to omit shadows from specular illumination at rendering time and still create an image with proper shadows. This will save computing time, if hard shadows are used, or memory, if soft shadows are produced. Note that this applies only to **RenderZone***. **Shaded Render*** does not generate specular reflections and, therefore, the **Render With Shadows** option in the **Shaded Render Options** dialog is disabled.

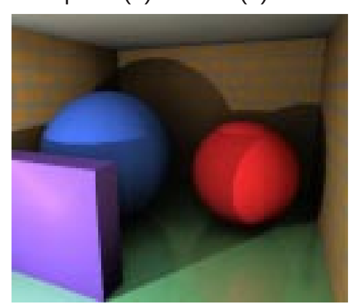

Figure 7.1.3.6: Diffuse shadows.

Figures 7.1.3.5 and 7.1.3.6 illustrate the effect of specular and diffuse shadows. The former Figure shows a sphere with a surface style that uses the **Plastic** reflection shader with a high specular reflection. The sphere is obstructed from the light by another object. In 7.1.3.5(a) the scene is rendered with both the **Calculate Shadows** option, in the **Radiosity Options** dialog, and the **Shadows** option, in the **RenderZone Options** dialog, on. The shadows are computed properly. In 7.1.3.5(b) **Shadows** is off (**RenderZone Options** dialog). The hotspot on the sphere is still visible, even though the sphere is in shadow.

Figure 7.1.3.6 shows the same scene, but a surface style with the **Matte** reflection shader, which has no specular parameter, is assigned to the sphere. Even though **Shadows** is off in the **RenderZone Options** dialog, the shadows are displayed correctly. This is because there is no specular reflection to cast shadows, which to be calculated correctly would have required the **Shadows** option to be on. By turning this option off, we are able to generate correct shadows, while saving memory and processing time.

Adding ambient for the reflected light

At any given time of the radiosity solution, only a percentage of the total amount of light has been assigned to the radiosity mesh polygons. Because of this and especially in the earlier stages, when rendering a radiosity solution, it is typically underexposed and appears dark. This may raise the need to temporarily add light, which is done with the following option.

Apply Reflected Light As Ambient: When this option is selected, the amount of light which has been emitted by all primary sources, but which has not yet been recorded at the mesh polygons is applied as a constant ambient light to all mesh polygons. During the early stages of a radiosity solution, this option will cause a scene to appear brighter than the final result. On the other hand, if this option is turned off, a scene only shows effect from light that has been recorded at the mesh polygons so far. For example, if a radiosity solution is 60% complete, only 60% of the light is rendered in scene, which consequently looks darker than the final result.

Figure 7.1.3.7 shows two images based on a solution which is about 60% complete. The first image is rendered with the **Apply Reflected Light As Ambient** option turned on, whereas the second image does not use this option.

a **b**

Figure 7.1.3.7: Two images rendered with the **Apply Reflected Light As Ambient** option (a) on and (b) off.

Current statistics

Given that the radiosity process is best executed in an iterative manner, it is frequently useful to know how far the radiosity solution has progressed at a given time. This is provided by the three statistics contained in the **Current Statistics** box. Note that these are simply informational fields and their content cannot be edited or changed.

% Of Light Absorbed: As a radiosity solution progresses, more and more light is "absorbed." That is, it is incorporated in the intensities attached to the radiosity mesh. This statistic displays the quantity of the absorbed light as a percentage of the total amount of light that would be absorbed by a 100% complete solution.

 # Of Iterations: The radiosity process is iterative and this statistic displays the number of iterations completed at the current stage of the radiosity solution.

 # Of Mesh Polygons: The radiosity solution is based on a progressive and recursive subdivision of all surfaces, resulting in the radiosity mesh. How many these are at a given time is displayed by this statistic.

These statistics can be read at any time the radiosity process is interrupted (stopped), but not exited. Recall, that a radiosity solution may stop by itself or it may be interrupted manually. When the **Radiosity Options** dialog is opened after such an interruption, the statistical values are shown in their respective fields. Before a solution is initialized, all these fields show zero values. It is possible to alter most of the parameters in the **Radiosity Options** dialog after a solution has been interrupted and even after it has reached its final goal. For example, if it turns out that after the initially specified 500 iterations the image generated from the solution is not satisfactory, it is possible to change the number of iterations to 1000 and let the solution continue.

Using preset radiosity parameters

A pop up menu at the lower left end of the **Radiosity Options** dialog (Figure 7.1.3.8) allows you to set the parameters in this dialog to preset values, which are organized at 10 levels.

Preset Parameters: Selecting an item from this pop up menu sets the parameters of the dialog to preset values. The higher the number chosen for the mesh, the better the quality will be. However, higher quality settings will also require longer times for generating the radiosity solution.

1 (Low Quality), 2, ..., 9, 10 (High Quality): Selecting one of these items sets the parameters of the dialog to a respective level of quality. These settings are summarized in the following table:

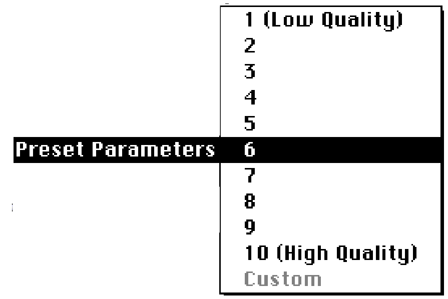

Figure 7.1.3.8:
The **Preset Parameters** pop up menu in the **Radiosity Options** dialog.

	1	2	3	4	5	6	7	8	9	10
Initialization										
Area Source Density	10	20	30	40	50	60	70	80	90	100
Secondary Source Density	10	20	30	40	50	60	70	80	90	100
Mesh Density	10	20	30	40	50	60	70	80	90	100
Resolve Intersections	–	–	–	–	–	–	–	–	√	√
Ignore O/F Meshing Attrs.	√	√	√	–	–	–	–	–	–	–
Sampling										
Min # Of Samples	1	1	1	1	1	3	3	3	3	3
Max # Of Samples	3	3	3	3	3	6	6	9	15	15
Quality	Simple	Simple	Simple	Better	Better	Better	Best	Best	Best	Best
Generation										
Adaptive Meshing	–	–	–	√	√	√	√	√	√	√
Density	–	–	–	55	60	65	70	75	80	85
Threshold	–	–	–	35	30	25	25	20	15	10
Continuous Meshing	–	–	–	–	–	–	–	–	√	√
Decrease Meshing	–	–	–	–	–	–	√	√	–	–
No Meshes After Processing Lights	–	–	–	√	√	√	–	–	–	–
Terminate After										
% Of Light Absorbed	70	70	70	80	80	80	90	90	90	90

7.1.4 Setting the radiosity meshing attributes

Radiosity meshing and its parameters can be assigned as an attribute to a complete object or a face of an object, using the Set Rendering Attributes tool. When radiosity meshing has been applied to individual entities (objects or faces) these can be used instead of the global meshing conditions specified in the **Radiosity Options** dialog, when a radiosity solution is derived.

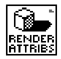

Set Rendering Attributes

This tool, located on the 12th row of the modeling tool palette, can be used to assign different rendering attributes to an object, including the radiosity meshing attribute. Which rendering attributes are assigned is determined by the options selected in the **Rendering Options** dialog (Figure 7.1.4.1). This dialog can be invoked directly from the tool.

To assign the radiosity meshing attribute to an object or face of an object, select the **Radiosity Meshing** option in the **Rendering Options** dialog and set the other parameters in the dialog. Then set the Topological level to Object or Face and with the Set Rendering Attributes tool active select an object or a face. To eliminate the radiosity meshing attribute from an object of face, deselect **Radiosity Meshing** in the dialog and select the respective entity as above.

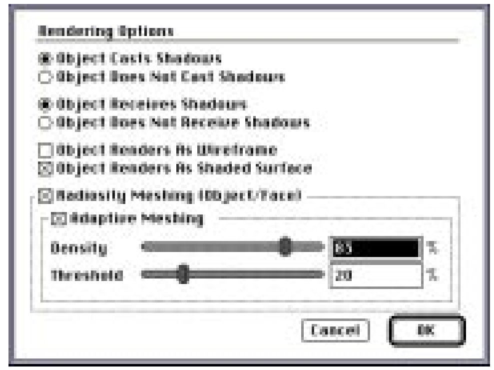

Figure 7.1.4.1: The **Rendering Options** dialog.

Note that **Radiosity Meshing** is the only rendering attribute in the **Rendering Options** dialog that can be assigned at the object or face level. All the others can only be assigned to objects.

When **Radiosity Meshing (Object/Face)** is selected in the **Rendering Options** dialog, **Adaptive Meshing** can also be selected. When it is on, it works as when set in the **Radiosity Options** dialog and is specified by the same parameters, namely, **Density** and **Threshold**. What is significant is that it is possible to completely exclude selected objects/faces from adaptive meshing, while it is applied to the rest of a scene, and vice versa. That is, the ability to apply radiosity meshing globally, individually, or by combining the two offers extensive flexibility in controlling the level of the radiosity solution. Recall that whether the individual meshing parameters are used depends on the status of the **Ignore Face/Object Level Meshing Attributes** in the **Radiosity Options** dialog.

7.1.5 Radiosity preview options

During the execution of a radiosity solution, it is important to receive visual feedback about the status of the current solution. The preview may be drawn in **Wire Frame**, **QuickDraw 3D** (Macintosh), **OpenGL** (Windows), or **Shaded Render**. Which type is used is determined by the currently selected display type in the **Display** menu. If any other display type is active when **Initialize Radiosity*** or **Generate Radiosity Solution*** is selected, it changes to wire frame. The preview is drawn in all open windows of the current project, allowing the preview of multiple views simultaneously. Whether the image is updated during the radiosity process is set in the **Radiosity Preview Options** dialog (Figure 7.1.5.1.). It is invoked from the **Preview Options...** button in the **Radiosity Options** dialog.

Update Display: When this option is on, a preview of the radiosity solution is produced as the radiosity progresses. How frequently it is updated is set by the following parameters.

Update Every *n* Seconds: This option causes the radiosity preview rendering ito be updated every ***n*** seconds, where ***n*** is entered in the field of this option.

Update After *n* Iterations: When this option is selected, the frequency of the updates is determined by the number of iterations, which is entered in the text field of this option.

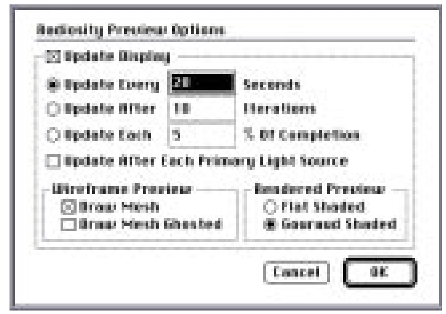

Figure 7.1.5.1: The **Radiosity Preview Options** dialog.

Update Each *n* % Of Completion: When this option is selected, a preview is drawn each time the specified amount of completion percentage has been reached.

Update After Each Primary Light Source: When this option is selected, a preview is drawn after each primary light source has been processed.

Generating a preview rendering takes time away from the computation of the radiosity solution. It is therefore not advisable to set the time below 10 seconds, since in such a case more time would be spent on drawing than generating the solution. When a Gouraud shaded rendering is used, it takes certain shortcuts in order to allow reasonably fast renderings. For example, textures, reflections, and transparencies are ignored.

Wire Frame Preview: These two options apply when **Wire Frame*** is used. **Draw Mesh** displays the radiosity mesh with the preview and **Draw Mesh Ghosted** draws the radiosity mesh in the project's ghost color, while the original objects are drawn in their colors.

Rendered Preview: These two mutually exclusive options apply when **Shaded Render*** or **QuickDraw 3D*** (Macintosh) is selected for the display. In these cases the preview can either be **Flat Shaded** or **Gouraud Shaded**. Note that, when **OpenGL*** is selected on a Windows machine, the preview display is always smoothly shaded.

7.1.6 Radiosity in form•Z Imager

The **RenderZone RadioZity** version of **form•Z Imager** can be used to generate radiosity based shaded render and **RenderZone** images. Since the **Imager** is not interactive, its primary use is for generating multiple high quality radiosity based images once the desired radiosity parameters have been established. For complete details on using the **Imager**, see section 3.9 of the **form•Z** User's Manual.

The radiosity based renderings are selected through two items at the end of the **Image Types** pop up menu in the **Imager Set** window (Figure 7.1.6.1) and in the **Default Image Type** pop up menu in the **Imager Set Up** dialog (Figure 7.1.6.2).

Shaded Render Radiosity executes a shaded rendering based on radiosity, and **RenderZone Radiosity** executes a **RenderZone** rendering based on radiosity. For both of these renderings the radiosity calculations are performed prior to the rendering using the radiosity options stored in the project. When there are sequential radiosity based renderings from the same project in the Imager set list, the radiosity calculations are only performed once and then used for all the renderings for that project.

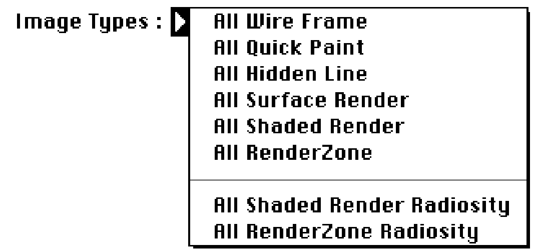

Figure 7.1.6.1: The **Image Types** pop up menu.

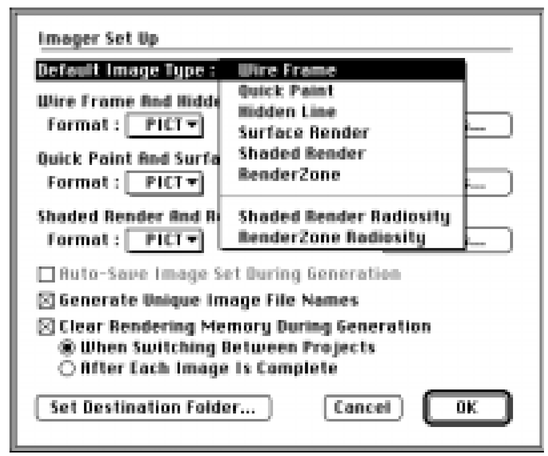

Figure 7.1.6.2: The **Imager Set Up** dialog and the **Default Image Type** pop up menu.

7.2 Lighting in radiosity

A major advantage of radiosity is its ability to simulate the distribution of light in a physical environment far more accurately than any of the other rendering methods that use simple illumination models. To achieve this, radiosity employs two additional types of lights that are not available to the other rendering modes: ***area*** and ***custom*** lights. In addition, the intensity for both these lights and the lights that are also available in the other rendering types are defined using accurate parameters borrowed from photometry. The most significant units of measurement used for lights are summarized below.

In Physics, a basic unit of measurement is ***power***. It represents energy per unit time interval. For light calculations, power is used to determine how much electric energy is used by a light source. Electric power is measured in ***watts*** (volt*ampere), which in itself gives no indication about how much light is emitted by a light source. To measure the light output, we also need to specify how much of the electric power supplied to a light is converted to visible radiation. This is referred to as the ***efficiency*** factor. A combination of this factor and watts is used to describe the output power of a light. For example, a 60 watt light bulb at 40% efficiency outputs 24 watts of light energy.

Luminous flux is the quantity of light energy emitted by a light source in a unit time interval. It is measured in ***lumens***. This measurement is used frequently by manufacturers of light fixtures and light bulbs as an indicator of the output power of a light. For example, a standard 60 watt incandescent light bulb has about 800 lumens. Unlike the energy specification of light sources in watts, lumens is a direct indicator of the brightness of a light source. That is, it does not represent the energy supplied to the light source, but describes how much energy is emitted.

Luminous intensity is the quantity that describes how much light energy is emitted by a light source in a unit time interval in a specific direction. It is measured in ***candelas***. While luminous flux indicates the overall light output, luminous intensity measures how strong a light is in a certain direction. This measurement also takes into account the volumetric shape of the light.

For example, a standard light bulb can be described as emitting light in a spherical fashion, with the light intensity being equal in all directions. To keep it simple, we will ignore the fact that a light bulb does not emit any light in the direction of its socket. If the luminous flux of this light is defined as 1000 lumens, the luminous intensity of the light can be calculated by dividing 1000 lumens by the volumetric angle of the light, in this case $4*\pi$ for the spherical shape of the light bulb. This yields a luminous intensity of about 80 candelas, which tells us how strong the bulb shines in any given direction. Next consider a reflector type light bulb with the same luminous flux. Its volumetric shape can be approximated by a conical sector of a sphere. Let us assume that the volumetric angle of this cone is one quarter of a sphere. The luminous intensity of this light source would yield 360 candelas. In other words, while the reflector light has the same luminous flux at 1000 lumens, its light energy is concentrated in a smaller area and therefore has more luminous intensity.

In **form•Z RenderZone RadioZity**, all three quantities can be used to specify the output energy of light sources.

Luminance: This quantity describes how much of the light energy that hits a surface is reflected back into the environment at a given direction per square area. Its units are measured in ***candelas*** per square foot or square meter.

Illuminance: This quantity describes the luminous flux (lumens) that hits a surface in a unit square area. Its units are ***lux,*** defined as one lumen per square meter, or ***footcandle***, which is defined as one lumen per square foot. Illuminance is a useful measure for determining the light intensity on a unit area of a surface. For example, when designing lighting for interior spaces, building code requirements specify the light intensity on a working desk in lux or footcandle, to ensure that enough light is present on the desk surface for proper working conditions.

In **form•Z RenderZone RadioZity**, when executing a **Shaded Render*** or **RenderZone*** rendering that is not based on a radiosity solution, the intensity of a light is determined by the value entered in the single **Basic Intensity** field, in the **Light Parameters** dialog. In contrast, there are four distinct dialogs for specifying the **Radiosity Intensity**: **Cone And Point**, **Distant**, **Area**, and **Custom**. They are invoked by clicking on the **Radiosity Intensity...** button found in the **Light Parameters** dialog. They are discussed in this section.

Note that, in radiosity, there is no projector light. When a projector light is used it is ignored by the radiosity solution. Also note that some of the radiosity lighting behavior is determined by parameters contained in the surface style assigned to a surface, rather than exclusively by the parameters of the lights. These are discussed in the following subsection.

7.2.1 Surface styles in radiosity based renderings

When the radiosity process computes the reflection of light off a surface, it is affected by two parameters contained in the definition of the surface style assigned to a surface: the ***color*** and the ***diffuse reflection*** factor.

Under certain illumination conditions, the color of a surface is reflected on neighboring surfaces. For example, a red table top in front of a white wall will reflect some of its red color on the wall, when illuminated from above. This effect is known as ***color bleeding*** and is illustrated in Figure 7.2.1.1. The color is contained in the surface style definition assigned to a surface.

Figure 7.2.1.1: An example of color bleeding.

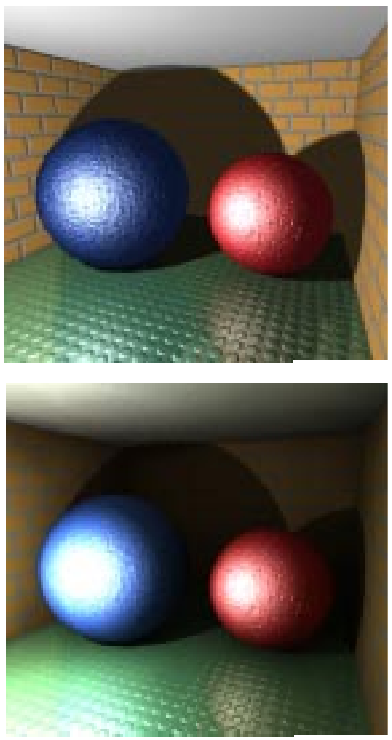

Figure 7.2.1.2: Displaying bumps (a) wthout and (b) with a radiosity solution.

A certain amount of light is diffusely reflected from a surface. This amount is defined by the **Diffuse Reflection** factor, present in the **Matte**, **Plastic**, **Chrome**, **Accurate Glass**, **Accurate Metal**, **Mirror**, and **Environment** reflection shaders. The **Constant**, **Metal**, and **Simple Glass** reflection shaders are assumed to have a diffuse reflection parameter of zero.

Radiosity assumes that the diffuse reflection is constant across a surface. That is, scaling the diffuse reflection via a reflection map is not taken into account during radiosity processing. As a general guideline, diffuse reflection values should be defined in a reasonable range to achieve realistic looking images. In the real world, diffuse reflection ranges from about 10% to 90% of the incident light. Diffuse reflection of 0%, where all light is absorbed, or 100%, where all light is reflected does not exist. Diffuse reflection values for some common surfaces are as follows:

Painted Ceiling	60 - 90%
Floor	15 - 35%
Aluminum	60 - 70%
Walls	50 - 80%
Wood Surfaces	30 - 50%

As a general guideline, a high diffuse factor should be used for a surface style, unless the surface is highly glossy, in which case a high specular reflection factor should be specified.

While diffuse reflection is calculated during the radiosity process, all other reflection parameters are calculated during a radiosity based rendering. In addition, the ambient reflection factor is completely ignored during such a rendering, since ambient reflection is now correctly calculated as illumination from secondary light sources. Since specular reflection is still computed at rendering time, the light intensity used in the specular calculation is taken from the simple illumination model. That is, the **Basic Intensity** value shown in the **Light Parameters** dialog is used for specular illumination.

In a radiosity model, rendering ***the effects of bump mapping are limited***. A rendering that uses the simple illumination model is able to simulate the surface irregularities, resulting from bump mapping. The bumps are created mostly from the diffuse reflection calculation at each pixel on the basis of an intentional disturbance of the surface normal direction at that pixel. These normal disturbances cannot be generated during a radiosity process, since radiosity simulates a process which accurately describes a physical phenomenon, whereas the normal disturbances are merely a trick to bend a surface at a pixel where in reality the surface is flat. In other words, the effects from the bump shaders used by a surface style will not be visible through diffuse reflection in radiosity based rendering. However, bumps will show up as a result of specular or mirrored reflections and refracted transparencies.

Figure 7.2.1.2 illustrates how bumps are displayed in radiosity based renderings. On the left is an image which is not based on a radiosity solution but was generated with the simple illumination. Bumps are visible on the wall surface, from diffuse reflection of the brick bumps. On the floor, the mirrored reflections are broken by the round bumps assigned to the floor's surface style. The bumps in the area of the hotspots of the spheres are generated by specular illumination. On the right is the same scene, but rendered after a radiosity solution was generated. Note that the bumps on the wall, which were generated previously through diffuse reflection, have disappeared, but that the bumps on the reflective floor and on the sphere are still visible.

7.2.2. Cone and point lights

Clicking on the **Radiosity Intensity...** button in the **Light Parameters** dialog while **Cone** or **Point** is selected from the **Type** menu invokes the **Radiosity Intensity: Cone And Point** dialog shown in Figure 7.2.2.1. In this dialog the intensity of a light is defined by using one of two mutually exclusive types of parameters: **Radiometric** or **Photometric**.

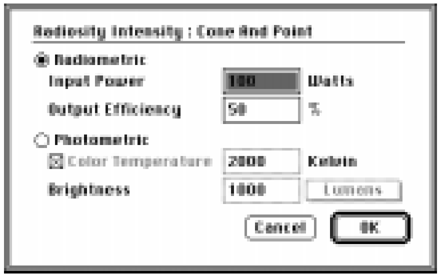

Figure 7.2.2.1: The **Radiosity Intensity: Cone And Point** dialog.

Radiometric: When this option is selected, the light's intensity is defined in ***watts***, using two parameters.

Input Power: This parameter value specifies the energy supplied to the source in **Watts**.

Output Efficiency: This parameter, expressed as a percentage, determines how efficient the light source is.

When multiplying the two parameters, the output energy of the light source can be calculated. It is common for light bulb manufacturers to specify the input energy of a light bulb. For example, a 60 Watt incandescent bulb consumes 60 Watts of electric energy, but outputs much less light energy. The efficiency of incandescent bulbs is around 40%, which means that a 60 Watt bulb creates about 24 Watts of light energy, with the rest being mostly converted to thermal energy.

Photometric: When this option is selected, the light energy is specified in **Lumens** or **Candelas**. Which unit is used can be selected from the pop up menu next to the text field. Recall that **Candelas** determines the light intensity at any given angle of the source. It is a unit which indicates the brightness of a light source at a given point. **Lumens** specifies the overall intensity of the light source and is a measure of the strength of the light. When specifying the light intensity in photometric units, it is also possible to determine the color temperature of the light.

Color Temperature: When this option is selected a value can be entered in the text field, which indicates how hot the light burns. This temperature, measured in degrees Kelvin, is converted into a color, which is mixed with the light's color selected in the **Light Parameters** dialog. Low color temperatures generate an orange color. The hotter the temperature, the closer to white the color of the light is.

Brightness: This value is expressed in Lumens and the higher it is the brighter the light is.

7.2.3 Distant or sun lights

When **Distant** light is selected in the **Type** menu, clicking on the **Radiosity Intensity...** button invokes the **Radiosity Intensity: Distant** dialog shown in Figure 7.2.3.1.

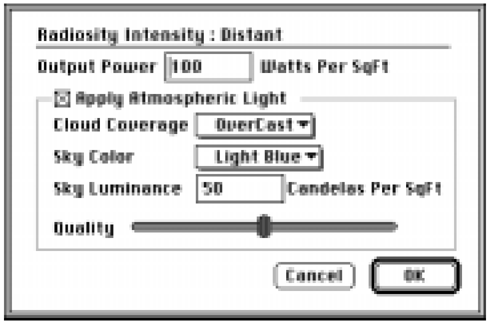

Figure 7.2.3.1: The **Radiosity Intensity: Distant** dialog.

Output Power: For distant lights, the intensity is not defined as an output quantity of the source, but as a quantity which indicates how much energy is received at a square unit on a surface perpendicular to the light. In most cases, distant lights are used to simulate the sun light. Instead of specifying the exact distance of the sun from the scene and the exact output power of the sun, it is more meaningful to specify the effect of this type of light on a surface.

Atmospheric lighting effects can also be set in the **Radiosity Intensity: Distant** dialog.

Apply Atmospheric Light: When this option is selected, indirect light which exists in the atmosphere, which for example reflects off the clouds, is added to the scene.

Figure 7.2.3.2: A scene rendered with a distant light using atmospheric effects.

Consider, for example, an interior scene of a room, with a window. Even if there is no light shining directly into the room, the environment still submits a considerable amount of light into the room. A scene rendered by a distant light, using atmospheric effects is shown in Figure 7.2.3.2. In this scene, the distant light does not shine directly into the room. Illumination is generated exclusively by the atmospheric light.

Cloud Coverage: From this menu, **Overcast**, **Mixed**, or **Clear** may be selected. **Overcast** generates a less uniformly directed light, whereas **Clear** orients the light more in the direction of the light source. **Mixed** generates an intermediate result.

Sky Color: From this menu **Light Blue**, **Medium Blue,** or **Dark Blue** can be selected, which determines the color of the sky. Note that the light emitted by the sky is not blue, but of varying shades of white.

Sky Luminance: This value indicates the brightness of the sky directly above in **Candelas** per square foot, or square meter in a metric project.

Quality: This slider bar controls the depth of the atmospheric light simulation. The higher the quality, the longer it takes to process this light during the execution of the radiosity solution.

7.2.4 Area lights

Cone, point and distant lights are idealized sources, assuming that the light originates from a single, infinitely small point. While this is adequate to simulate, for example, an incandescent light bulb, other light fixtures, such as neon tubes are hard to recreate with these light types. An area light is associated with a physical object, whose surfaces are emitting light. For example, it is possible to model a curved neon tube, using the sweep tool and a circular source shape, and turn the resulting object into an area light. This light source behaves exactly like a glowing neon tube. A scene rendered with such a light is shown in Figure 7.2.4.1.

Figure 7.2.4.1: A scene rendered with an area light.

Creating an area light

You can create an area light as follows:

- Create an object which will become an area light. That is, this object represents the shape of the area light.
- Preselect the object with the Pick tool.
- Create a new light. Recall that you do this by either clicking in the Lights palette under its last entry, or by clicking on the **New...** button in the **Lights** dialog.
- When the **Light Parameters** dialog is invoked, select **Area** from the **Type** pop up menu. This will cause additional **Area Light** options to be displayed at the lower end of the dialog, as shown in Figure 7.2.4.2.

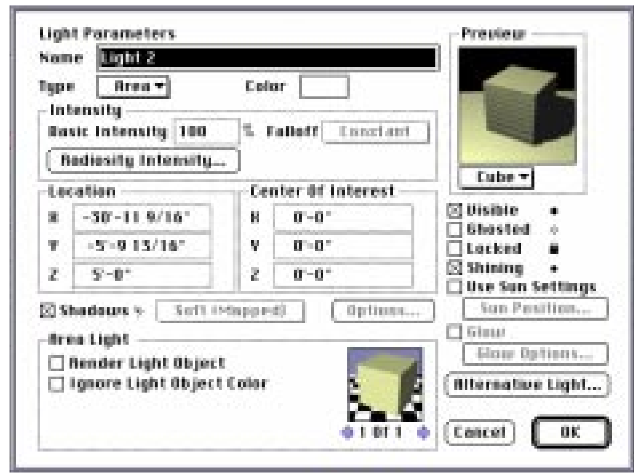

Figure 7.2.4.2: The **Light Parameters** dialog when an area light is selected.

- Select the desired settings in the **Area Light** group of options, whose functionality is discussed below.
- Click on the **Radiosity Intensity...** button to invoke the **Radiosity Intensity: Area** dialog, where proper intensity parameters are set, as discussed below.
- Set additional parameters in the **Light Parameters** dialog if desired and click on **OK**.

This completes the definition of an area light, whose name is now displayed in the Lights palette. To edit and change its parameters, double click on its name in the Lights palette. Any previously defined light of a different type can also be changed to an area light by double clicking on its name, provided an object is also preselected when editing the light. The object used for an area light can be a solid, a surface object, or a surface solid (double-sided surface). It cannot be an open vector line, a symbol, or plain text.

Once an object has been converted into an area light, it is automatically deleted from the Objects palette. That is, no more modeling operations can be performed on the object defining the area light. In order to retrieve the object, the area light must be converted to another light type. This can be done by executing the following steps.

• Double click on the area light's name in the Lights palette to invoke the **Light Parameters** dialog.
• In this dialog, select a type other than **Area** from the **Type** menu.
• After clicking **OK**, the new light type is drawn in place of the area light, and the object reappears as a regular **form•Z** object. The object preserves its original color, but is placed on the active layer.

Surfaces of an area light shine only in the positive direction, as indicated by the direction arrow. For example, if a surface object, which consists of a single face, is defined as an area light, the part of the scene which is behind the surface is not illuminated directly by this light.

Setting the intensity of an area light

The intensity of an area light is set through the **Radiosity Intensity: Area** dialog (Figure 7.2.4.3) invoked by clicking on the **Radiosity Intensity...** button, in the **Light Parameters** dialog.

Radiometric: When this option is selected, the total **Output Power** of the area light is defined in **Watts**.

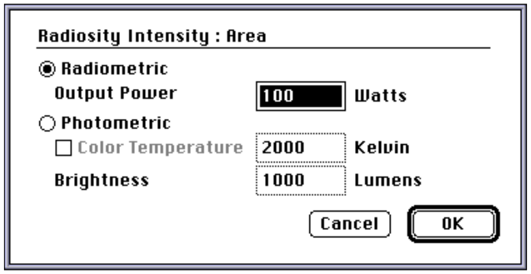

Figure 7.2.4.3: The **Radiosity Intensity: Area** dialog.

Photometric: When this option is selected, the overall light intensity is specified in **Lumens**. The **Color Temperature** option is as with cone and point lights.

Unlike cone and point lights, the photometric units of area lights can only be specified in Lumens. Candela values are only meaningful if the light source has an equal intensity in all directions. This is not the case with area lights, since the direction and size of the faces determines the intensity of the light at a given point in space.

Shadows

As with all lights, the area lights can cast shadows. However these are calculated differently than other lights.

Shadows: When this option is selected, the area light will cast shadows. Unlike direct, cone and point lights there are no parameters for **Soft** or **Hard** shadows, since radiosity does not use shadow maps or raytraced shadows. In other words, all the other parameters associated with the **Shadows** are dimmed and inactive.

Setting alternative lights

Because area lights are not available outside radiosity, when a scene that contains them is rendered using **Shaded Render*** of **RenderZone***, the option to approximate them with point lights is offered in the respective dialogs. The following option allows you to determine how area lights will be approximated outside radiosity.

Alternative Light...: Clicking on this button invokes the **Alternative Light Parameters** dialog shown in Figure 7.2.4.4. It can be used to set the parameters for an alternative point light.

The parameters which the area light and the alternative light have in common, such as the location, are presented in the **Light Parameters** dialog. The parameters which are unique to the point light are edited in the **Alternative Light Parameters** dialog. Note that the dialog includes shadow parameters.

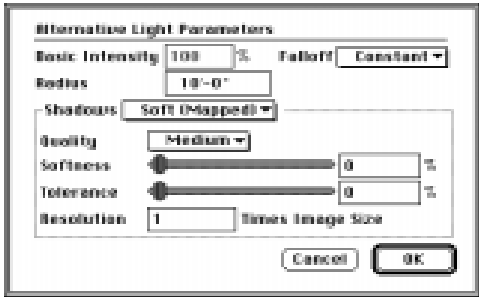

Figure 7.2.4.4: The **Alternative Light Parameters** dialog.

Area light parameters

When **Area** is selected from the **Type** pop up menu the **Area Light** box appears at the lower end of the **Light Parameters** dialog, where two parameters can be set and the surface styles assigned to the area light object may be inspected and possibly redefined.

Render Light Object: When this option is selected, the object of the area light is rendered in the image derived from the radiosity solution. Otherwise, the object is invisible. However, even if the object is rendered invisible, the energy emitted by the area light is still processed.

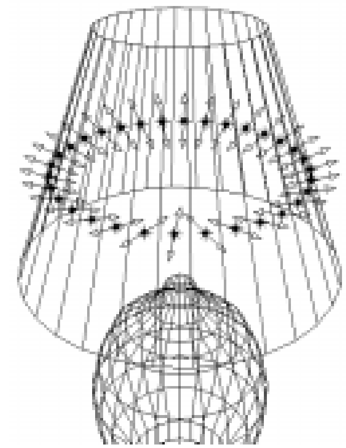

Figure 7.2.4.5: An area light, drawn in wireframe display.

Recall that, in wire frame, an area light and its object are displayed, when its visibility attribute in the Lights palette is on. In addition, its directions may be also displayed, if the **Show Directions** option in the **Wire Frame Options** dialog is on. The direction of the light leaving all surfaces is indicated by an arrow starting at the center of each face and pointing outward. An area light drawn in wire frame is shown in Figure 7.2.4.5.

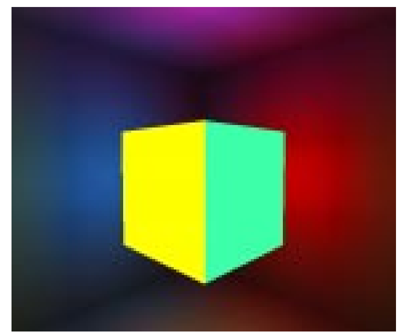

Figure 7.2.4.6: Illumination with an area light with different face colors.

Ignore Light Object Color: When this option is selected, the area light emits light in the color shown in the color box at the top of the dialog. When deselected, the light color is filtered through the color of the surface styles used by each face of the light object. This is illustrated in Figure 7.2.4.6.

A cube, which has different colors on each of its six faces, will, once converted into an area light with a white color, illuminate the scene with six different colors. Each color is most intense in the direction of the face, to which the color was assigned. This effect is shown in Figure 7.2.4.6. Note that this color is always a solid uniform color, regardless of the color shader of the surface style used by the faces. For example, if a texture map is used as the color shader, an average color is determined from the texture. In other words, the area light does not project the color pattern generated by the color shader onto the scene, but always uses a plain color.

7.2.5 Custom lights

In conventional computer renderings, lights are assumed to be distributing their rays uniformly. For example, a point light, which is typically used to represent a bulb hanging in a room is assumed to emit rays of equal intensity in all directions. However, light sources do not really work this way in the real world. For example, a bulb has a higher intensity in the direction of the socket, and its intensity falls off toward the socket, and it is zero directly behind the socket. The radiosity procedures are specifically intended to take into consideration these real world "irregularities," which is done through the use of a light type called ***custom light***. The custom light is essentially a method which allows us to represent variable intensities in different directions about a light source.

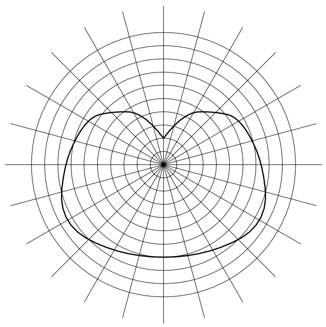

Figure 7.2.5.1: Light distribution curves.

The custom light representation

The varying intensities of a custom light can be visualized and represented as one or more curves called ***distribution curves***. Such a curve is drawn on a plane relative to a radial grid, as shown in Figure 7.2.5.1. The curve is defined by a set of points, each of which is defined by its ***angle*** from the horizontal orientation of the radial grid and its ***distance*** from the center of the grid. The distance represents the intensity of the light at the direction where the point is positioned. A custom light is completely defined by a set of planes, that intersect on the axis of the light. The axis is defined by the ***origin*** and the ***center of interest*** of the light. This is illustrated in Figure 7.2.5.2.

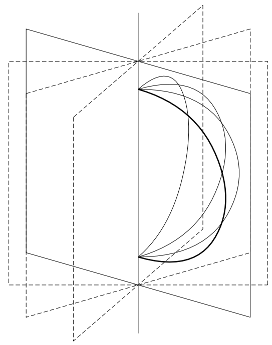

Figure 7.2.5.2: Light distribution curves in different planes.

The distribution curves and their planes are displayed when the intensities of a custom light are defined. When a custom light is displayed on the screen, it is represented as shown in Figure 7.2.5.3. It is represented as a sphere, by three circles that lie on perpendicular planes. Two of these planes intersect at the axis of the light, which has a direction determined by the origin and center of interest of the light. These points and their direction are represented by a line that connects then and an arrow head at the position of the center of interest. Another line that starts at the origin and is perpendicular to the axis, represents the zero angle position relative to which the angles of the planes used for the representation of the light intensities are measured. The zero angle line also ends at an arrow head which is smaller than the direction's arrow head.

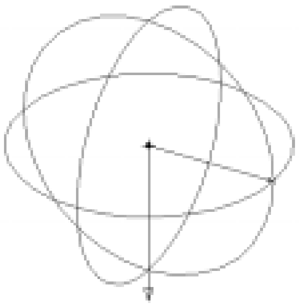

Figure 7.2.5.3: Wireframe display of a custom light.

Creating a custom light and setting its parameters

A custom light is generated as with all the other lights (except the area light). Clicking under the last entry in the Lights palette or clicking on the **New...** button in the **Lights** dialog invokes the **Light Parameters** dialog. In this dialog, select **Custom** from the **Type** pop up menu, which causes custom light specific options to appear in the dialog (Figure 7.2.5.4).

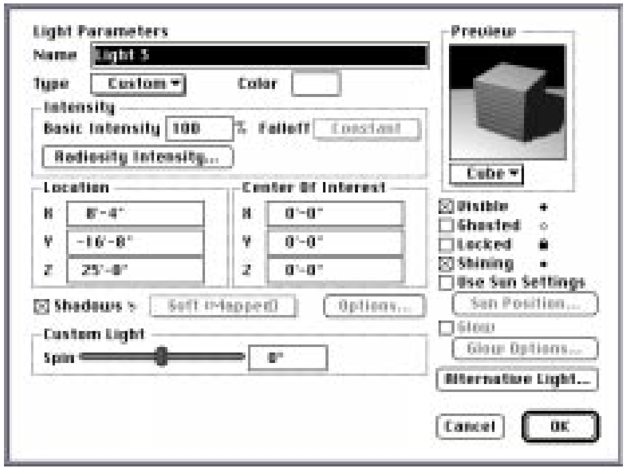

Figure 7.2.5.4: The **Light Parameters** dialog when a custom light is selected.

Shadows: As for the area light, this option can simply be turned on or off and there are no other options to be set. The shadow options available for other types of light are dimmed and inactive.

Spin: The only parameter that appears in the **Custom Light** box at the bottom of the dialog, this slider bar is used to set the spin angle, which is also displayed in the edit field next to it. The spin angle represents a rotation of the light about its axis and is measured relative to the zero angle line.

Alternative Light...: As with the area light, when a scene that contains custom lights is rendered without radiosity, the custom lights are treated as point lights whose parameters are specified by clicking on this button. If the parameters for the alternative light are not specified, default values are used.

Setting the intensities of a custom light

Clicking on the **Radiosity Intensity...** button while **Custom** is selected from the **Type** pop up menu in the **Light Parameters** dialog invokes the **Radiosity Intensity: Custom** dialog shown in Figure 7.2.5.5. This dialog is used to display a set of distribution curves representing the intensities of the custom light in different directions.

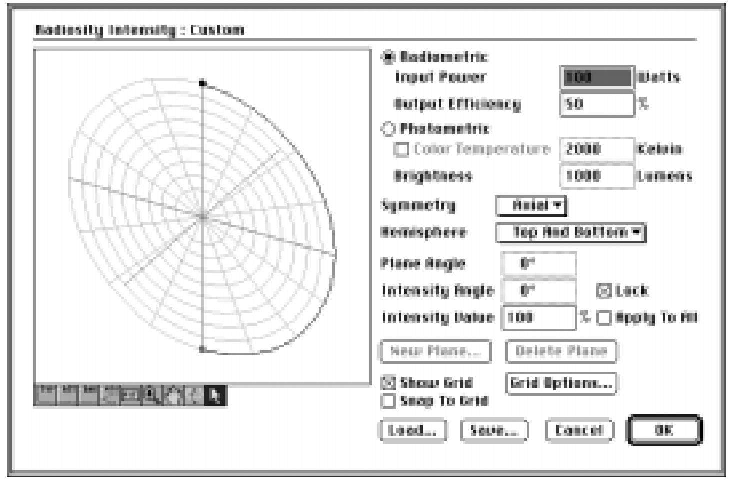

Figure 7.2.5.5: The **Radiosity Intensity: Custom** dialog.

The **Radiosity Intensity: Custom** dialog is a preview dialog. Its left portion contains a graphic area that displays the distribution curves of the custom light. A set of default curves are initially generated by the program, which vary depending on the selections from the **Symmetry** and **Hemisphere** pop up menus (see below). The distribution curves can be edited and changed according to the desired intensities.

The following conventions are used when displaying the distribution curves in the preview area:

- One of the distribution curves is the ***active curve*** and is displayed in black. There can be only one active curve at a time. When the **Show Grid** option is on, the grid of the active curve is also displayed.
- When there is more than one distribution curve, the inactive curves are displayed in the ghost color of the project (default is gray).
- One of the points of the active distribution curve is the ***active point*** and is displayed in the highlight color of the project (default is red).
- All the points of a curve that are not active are displayed in black.

The tools under the preview area can be used to manipulate the displayed view as in the other preview displays (see section 3.2.6 in the **form•Z** User's Manual).

The Arrow tool can be used to execute the following graphic operations:

• Clicking on a distribution curve makes it active. If the **Show Grid** option is on, the displayed grid switches to that of the new active curve.
• Clicking on a point of the active distribution curve makes it the active point.
• Clicking on the active curve while pressing the ***option*** (Macintosh) or ***shft+ctrl*** (Windows) keys inserts a new point at the position of the click.
• Clicking on a point while pressing the ***control*** (both platforms) key deletes the point. However, note that the endpoints of a distribution curve cannot be deleted.
• Clicking on a point and dragging it moves it along the plane of the distribution curve.

The active curve and active point affect the values displayed in the fields on the right portion of the **Radiosity Intensity: Custom** dialog. They are also the entities to which changes made through numeric input are applied. The options to the right of the dialog are as follows:

Radiometric, Photometric: These two mutually exclusive groups of options located at the top of the right side are the same as the options in the **Radiosity Intensity: Cone And Point** dialog and work as for the cone and point lights (see section 7.2.2). They are used to specify the overall intensity of the light, using either radiometric or photometric units.

Symmetry: One of four distinct symmetry patterns is selected from this pop up menu. The symmetry pattern establishes how the distribution curve planes will be used to completely describe the intensities of a custom light. Each symmetry pattern has its own minimum plane requirements, thus, each time a different symmetry pattern is selected a warning message is issued. After responding **OK**, the program proceeds and generates the default set of planes for the particular symmetry, as follows:

None: When this item is selected, which does not actually involve any symmetry, the light distribution requires at least three distribution curves. The first must be on a plane at 0 degrees. The second can be at any angle. The third must be at a plane angle larger than 180 degrees. When first selected, four planes are created automatically, at 0, 90, 180 and 270 degrees around the light direction.

Axial: When this item is selected, the light distribution requires only one distribution curve. The angle of the plane of the curve is always 0 and cannot be edited. This option produces a description of the light which is of a rotational symmetry. In other words this is equivalent to an object of revolution.

Quadrant: When this item is selected, the light distribution is symmetrical in each quadrant and requires at least two distribution curves, one at 0 and the other at 90 degrees. When first selected, curves at these two plane angles are created automatically.

Plane: When this item is selected, the light distribution is symmetrical relative to the 0 - 180 degree plane and requires at least two distribution curves, one at 0 and the other at 180 degrees. When first selected, curves at these two plane angles are created automatically.

Light distribution curves for all four symmetry options are illustrated in Figure 7.2.5.6.

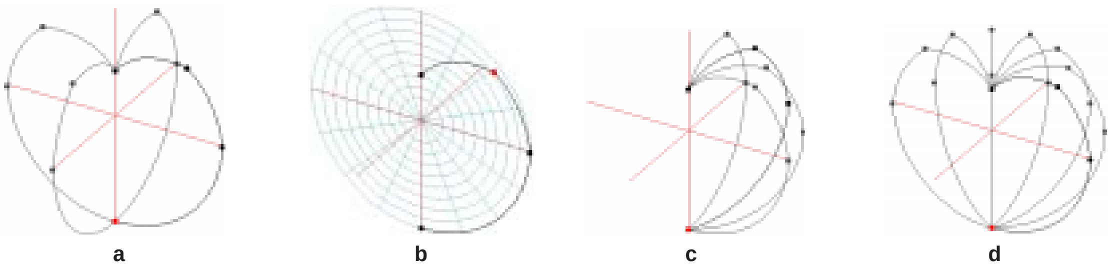

Figure 7.2.5.6: Distribution curves for the (a) **None**, (b) **Axial**, (c) **Quadrant**, and (d) **Plane** symmetry options.

Hemisphere: One of three mutually exclusive options is selected from this pop up menu. The selection determines the extent of the light relative to its hemispheric plane. This is the plane that is perpendicular to the axis of the light and passes through its origin. It splits the sphere of a light into two hemispheres.

Bottom, Top, Top And Bottom: Selecting one of these items causes the custom light to shine only from its bottom, top, or all around, respectively. These variations are illustrated in Figure 7.2.5.7. Note that even a hemisphere which generally shines may contain points of zero intensity that do not shine. These areas can be established by editing and manipulating the default distribution curves that are initially generated according to the selection from **Hemisphere**.

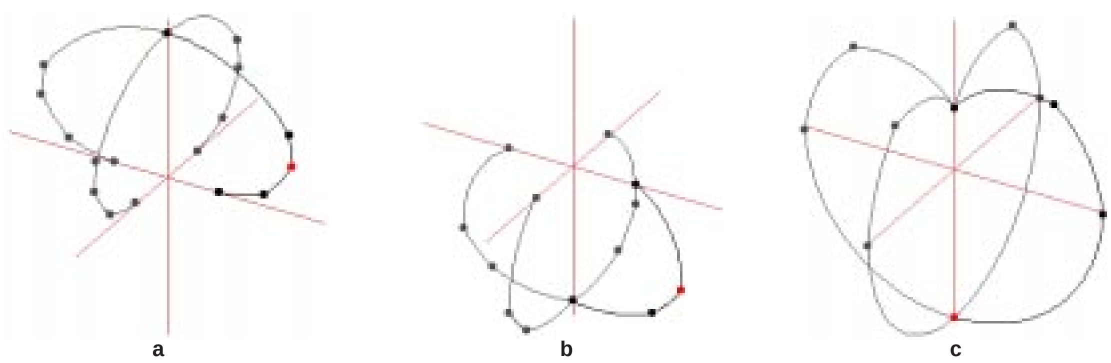

Figure 7.2.5.7: Distribution curves located in (a) the **Top**, (b) the **Bottom**, and (c) the **Top And Bottom** hemispheres.

Plane Angle: This editable field displays the angle of the plane of the ***active distribution curve*** relative to the zero angle plane. Changing its value repositions the plane of the active distribution curve. If the angle of this plane cannot be changed (as for the single distribution curve used by the **Axial** symmetry), this field is dimmed and inactive.

Intensity Angle: This editable field displays the angle of the ***active point*** (on the active distribution curve) from the origin of the light. Changing this value repositions the point. If the angle of an active point cannot be changed, this field is dimmed. This occurs when one of the end points of the distribution curve is active, or the **Lock** option next to this field is on.

Intensity Value: This editable field displays the intensity of the light at the position of the ***active point***. It is expressed as a percentage of the overall intensity defined by the radiometric or photometric units at the top of the dialog. In the preview diagram of the custom light the intensity is represented as a relative distance from the origin of the light.

Apply To All: When this option is off, which is the default, any change made to the **Intensity Value** will be applied to the active distribution curve only. If this option is on, the change will be applied to all the distribution curves.

New Plane...: Clicking on this button invokes the **New Plane Angle** dialog (Figure 7.2.5.8) where the angle at which a new distribution curve plane will be created can be entered. The dialog also includes a message informing what the range of acceptable angles is, which depends on the type of symmetry that is currently used. These conditions are summarized in the table of Figure 7.2.5.9. When you click on **OK**, a plane is created, a copy of the active distribution curve is made on it, and this new curve becomes the active distribution curve.

Delete Plane: Clicking on this button deletes the active distribution curve and its plane. If the active curve is one that cannot be deleted, this button is dimmed and inactive.

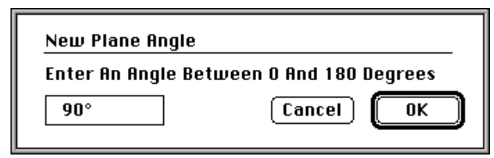

Figure 7.2.5.8: The **New Plane Angle** dialog.

Symmetry	Acceptable angle for new plane
None	0° - 360°
Axial	no new plane can be created
Quadrant	0° - 90°
Plane	0° - 180°

Figure 7.2.5.9: Acceptable angles for new planes per type of symmetry.

Show Grid: When this option is on, the radial reference grid of the active distribution curve is also displayed. When off, the grid is hidden.

Grid Options...: Clicking on this button invokes the **Radial Grid Options** dialog (Figure 7.2.5.10) where the parameters of the grid can be set. These parameters are the **Angle Spacing**, expressed in degrees, and the **Distance Spacing**, expressed as a percentage.

Figure 7.2.5.10: The **Radial Grid Options** dialog.

Snap To Grid: When this option is on and the position of the active point is graphically manipulated using the Arrow tool under the preview area, it can only be moved to the points of the grid. These are the points where the concentric circles intersect with the radii of the grid.

Load...: When this option is selected, a data file containing the parameters for a custom light may be loaded. This file may be in one of three industry standard formats.

- IES: Illumination Engineering Society of North America
- CIE: International Commission on Illumination
- CIBSE: Chartered Institution of Building Services Engineers (British)

These data files contain angles and values for one or more distribution curves as well as optional parameters for the output energy of the light. If energy output parameters are supplied, they will be entered in the respective fields. If not supplied, the parameters currently present are maintained. Most light fixture manufacturers supply data files in one of those three formats with their products.

Save...: When this option is selected, the current description of light distribution can be saved to a data file, formatted in one of the three standards.

7.2.6 Specifying the correct light intensity

The light intensity of all light types except distance lights can be specified in either lumens or watts. While the photometric intensity in lumens is the preferred unit with respect to generating accurate lighting results, in many cases the lumens rating of a bulb may not be known, while the electric input power in watts is. Unfortunately, there is no direct conversion from watts to lumens, since watts represents the electric input power and the light's output also depends on the efficiency with which electric energy is converted into visible light energy. The following charts contain the watts/lumens correspondence for various light types.

Incandescent

Watts:	*Lumens*:
25	230
40	430
60	730
75	960
100	1380
150	2220
200	3150
300	5000
500	8400
1000	18800

Low voltage halogen

Watts:	*Lumens:*
10	140
20	300
50	850
100	2000

High pressure mercury

Watts:	*Lumens:*
50	2000
80	3850
125	6500
250	14000
400	24000

High pressure sodium

Watts:	*Lumens:*
150	14000
210	18000
250	25500
400	48000
1000	120000

Fluorescent tubes

Watts:	*Lumens:*
20	800 - 1300
40	1900 - 3400
65	3100 - 5500

Note that the lumens rating of fluorescent tubes varies for the same input power because of differences in the color spectrum. Typical color temperatures for some standard lamp types are listed in the chart below.

Lamp type	***Color temperature in degrees Kelvin***
warm white fluorescent	3020
white fluorescent	3450
cool white fluorescent	4250
daylight fluorescent	6250
clear mercury	5710
improved color mercury	4430
low pressure sodium	1740
high pressure sodium	2100
tungsten halogen	3190

7.3 Radiosity tips

Radiosity is a complex process which typically requires fine tuning to produce desirable results. This section discusses a few techniques that will help you avoid some of the typical radiosity problems.

7.3.1 How to improve the radiosity solution

While "perfect" radiosity solutions are always possible by increasing the overall resolution of the radiosity mesh, this is typically not feasible, due to the processing times that extreme radiosity solutions would require. There are also techniques that can be employed to achieve desirable results without resorting to global resolution increases, as well as techniques for controlling the time required to complete a radiosity solution. Some of these techniques are discussed in the remainder of this section.

The radiosity solution takes too long

Despite all efforts to make radiosity as fast as possible, a radiosity solution may take a considerable amount of time to generate a satisfactory result, especially if a high image quality is required. This coupled with the iterative nature of the execution of the radiosity solution, where the calculations are repeated after adjustments are made to the parameters, may result in very lengthy computation times if not properly planned. The proper method here is to keep each stage of the iteration as simple as possible by only generating as much detail as necessary for evaluating a specific aspect of the radiosity solution.

When starting to generate a new radiosity solution, it is advisable to set the radiosity parameters in such a way that initially the solution progresses fast, but generates low quality images. While these images might not be usable as a final result, they are generally sufficient to tell whether the scene has been set up properly. For example, one will be able to examine whether the light sources are strong enough or whether the desired contrast in the scene has been achieved. As changes are made to the scene and the results become more and more satisfactory, it is possible to increase the quality of the radiosity solution, at the expense of a longer processing time.

One of the secrets to a fast radiosity solution is to keep the density of the mesh as low as possible and still generate a radiosity based rendering of a desired quality. A number of tricks can be applied to achieve this and are discussed in the following subsections.

Crisp shadows by manual meshing

Crisp shadow boundaries can always be achieved by increasing the resolution of the radiosity mesh. However, given than high resolution meshes require extensive time to compute, techniques other that simply increasing the resolution are more efficient. If a shadow boundary is a straight line, it helps to align the direction of the radiosity mesh with the direction of the shadow boundary. This can be achieved by applying **form•Z**'s Mesh tool to selected faces in a way such that the meshing direction and the direction of the shadow boundary match. Figure 7.3.1.1 shows such an example. Without applying the Mesh tool, the radiosity mesh has an arbitrary alignment and the shadow boundaries in Figure 7.3.1.1 (a) show artifacts. In Figure 7.3.1.1 (b), the floor on which the table stands has been subdivided with the Mesh tool. The direction of the mesh and the edge of the table align, which allows the radiosity mesh to follow the shadow boundary projected by the table top. In both cases the number of mesh polygons is about the same. Note that the second image shows no artifacts at the shadow boundaries.

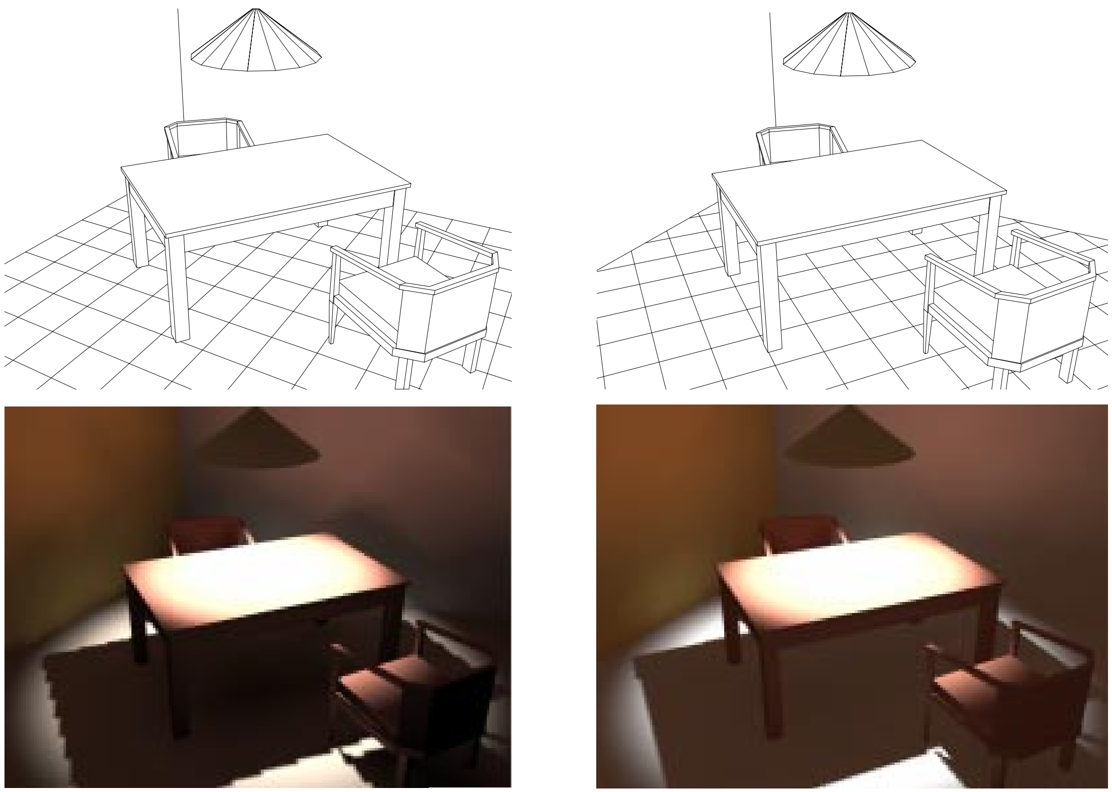

a **b**

Figure 7.3.1.1: Rendering shadows (a) without and (b) with aligning the radiosity mesh with the shadow boundary.

The next example is a case where we have no opportunity to align the direction of the mesh and the shadow boundaries since the latter has a curved shape. Consequently, other methods are applied.

Rough shadow boundaries are caused by a coarse radiosity mesh, combined with hard shadow edges. Such a situation is shown in Figure 7.3.1.2 (a) where a cone light is shining on a corner of a room. Increasing the radiosity mesh density or softening the shadow boundaries will reduce or eliminate the problem. Increasing the mesh density can be done on a global level by increasing the meshing parameters in the **Radiosity Options** dialog, or on a local level, by assigning meshing attributes to individual objects or faces by applying the Set Rendering Attributes tool. A denser mesh is illustrated in Figure 7.3.1.2 (b). Alternatively, the shadow boundaries may be softened without having to increase the mesh density. When using a cone light this can be done by assigning a smaller inner cone angle. Such a solution is shown in Figure 7.3.1.2 (c). Using area lights will also create softer shadows. More examples where rough shadow boundaries are avoided by proper selection of light type are discussed in the next section.

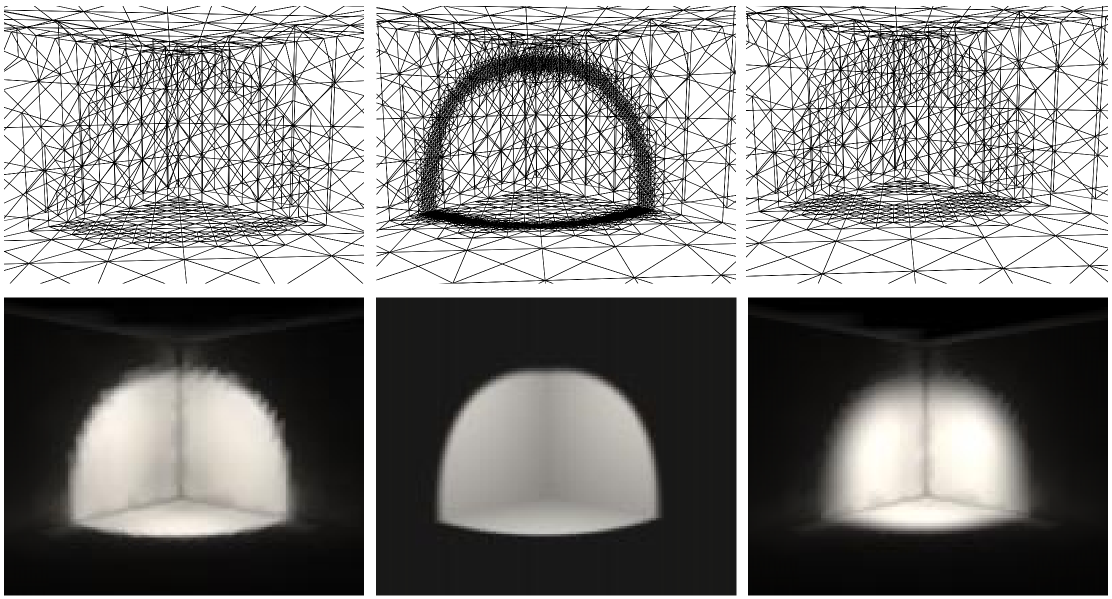

a **b** **c**

Figure 7.3.1.2: (a) Rough shadow boundaries; (b) denser meshing leading to crisper shadow edges; (c) softer shadows from a smaller inner cone angle.

Crisp shadows by proper choice of light type

Another way to keep the mesh density down is to avoid hard shadows whenever possible, which can be done by choosing the proper light type. Consider a table lamp, with a cylindrical lamp shade. One way to model this lamp would be to put a point light in the center of the lamp and let the lamp shade generate the shadow it casts on the scene. The shadow boundaries created by the top and bottom rim of the lamp shade would require a dense mesh to avoid artifacts. A rendering with such an approach is shown in Figure 7.3.1.3 (a).

A better solution would be to put two cone lights inside the lamp, one shining straight up and the other shining straight down. Cone lights allow for a soft edge, by setting the inner cone angle a few degrees smaller than the outer cone angle. If the cone lights are lined up in such a way that they do not intersect the rims of the lamp shade, no shadows need to be generated. In addition the soft edge of the cone light will allow for a lower density mesh without visual artifacts. A rendering with this setup is shown in Figure 7.3.1.3 (b).

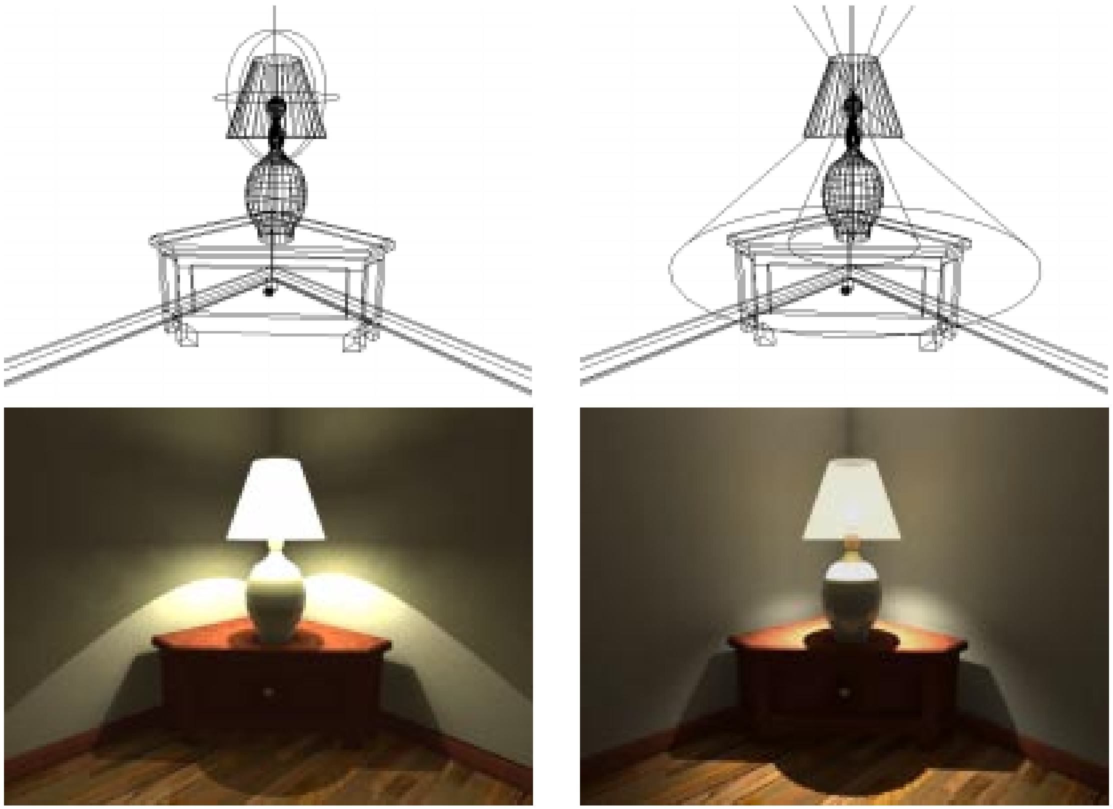

a b

Figure 7.3.1.3: Simulating soft shadows by using (a) a point light and (b) two cone lights.

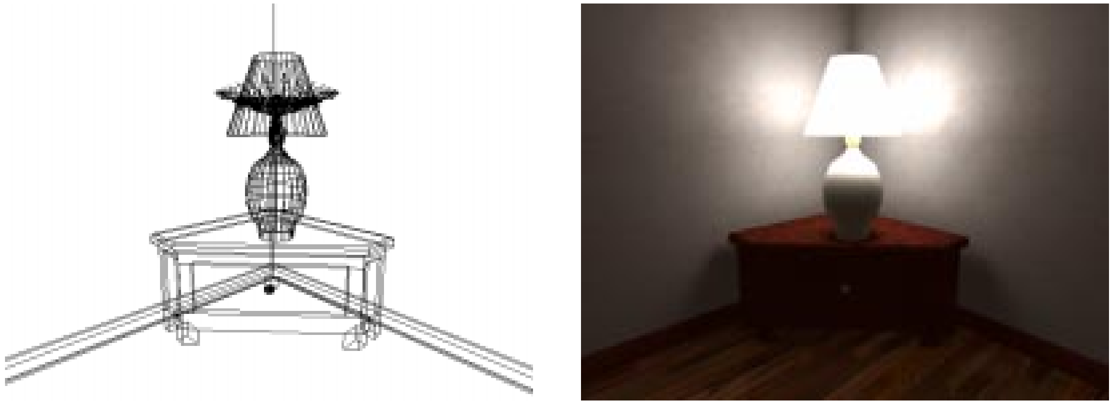

Figure 7.3.1.4: Simulating soft shadows by using area lights.

A third method to avoid hard shadow edges would be to make use of area lights. Unlike all other light types, where light is emitted from a single point, area lights emit light from points randomly sampled across the faces of the area light object. As a result, the shadows caused by the light emitted from one point on a face may be softened by the light emitted by another point of the area light. Figure 7.3.1.4 shows the same lamp as above, except this time the lamp shade is converted into an area light. The image generated from the radiosity solution shows no visual artifacts, although the mesh density is the same as in the examples above.

Area lights may also be used efficiently to simulate the illumination of a room by an exterior window, if no direct sun light shines through the window. In this case, the area light represents the non directional light existing in the outside environment, which radiates non uniformly through the window into the room. Such an example is illustrated below in Figure 7.3.1.5.

Figure 7.3.1.5: An area light used as an exterior window.

7.3.2 Troubleshooting

Below is a list of problems which may arise and suggestions how to deal with them.

Shadow or light leaks

A common problem of radiosity based renderings are shadow or light leaks. Such artifacts are illustrated in Figure 7.3.2.1 (a). Note that at the area where the leg of the stand touches the floor, a dark shadow band is displayed.

When looking at the radiosity mesh of this scene, one can observe that the mesh polygons of the floor where the leg touches the floor are partially under the leg and partially visible. When light intensities are computed, one portion of those mesh polygons receives no light, whereas the other portion is illuminated. During the rendering of the scene, the light intensities are interpolated across the mesh polygons. Interpolating the dark color that is under the leg with the light color which is visible, at the edge of the leg produces a shade that appears to be a leak.

The solution to this problem is to insert edges where objects touch. One way to achieve this in the example would be to union the leg with the floor. Alternatively, the **Resolve Intersections** option in the **Radiosity Options** dialog may be selected, which will automatically split intersecting objects. After this has been done, the mesh polygons are not partially obscured anymore and the resulting rendering does not show any more shadow leaks, as illustrated in Figure 7.3.2.1 (b).

a

b

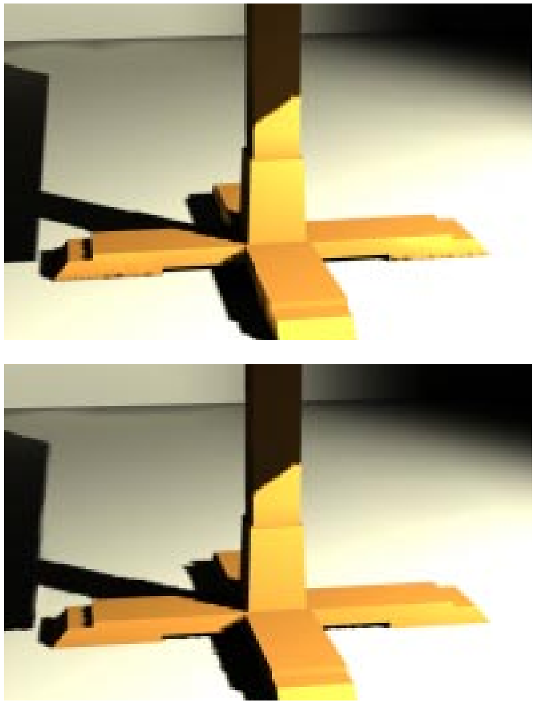

Figure 7.3.2.1: (a) Shadow leaks and (b) corrected leaks.

The radiosity mesh becomes too dense

It is possible that a radiosity solution very quickly generates a large number of mesh polygons. This in return will have a strong impact on the time it will take to arrive at an accurate solution. Dense areas of the radiosity mesh are necessary in regions, which have sharp contrast in illumination, such as hard shadow boundaries. The answer to excessive meshing is to avoid over-meshing areas which do not exhibit sharp contrast, or which are too dark to see. A number of steps may be taken to achieve this. First the **Threshold** parameter in the **Radiosity Options** dialog for global meshing or in the **Rendering Options** for object/face level meshing may be increased. Large values for this parameter will cause less subdivision. Figure 7.3.2.2 (a) shows a scene with a dense mesh created with a **Threshold** parameter of 10%. Note that the mesh is dense even in areas where there is only little light contrast, such as in the center of the illuminated arcs. The same scene has been rendered based on a solution processed with a **Threshold** parameter of 25%. Note that the Gouraud shaded renderings are virtually the same. However, the mesh of the second example contains only a fraction of the number of polygons of the first example.

A second solution to overmeshing would be to assign a denser mesh only to objects or faces which contain areas of sharp contrast, but to process the rest of the scene with a lower mesh. This can be done, by applying the Set Rendering Attributes tool to those objects or faces and by setting stricter meshing parameter in the **Rendering Options** dialog.

A third solution would be to generate a view dependent radiosity mesh. If renderings based on such a mesh are always created form the same view, areas in the scene which are not visible can be safely skipped from meshing. This can reduce the mesh density by a large amount, resulting in a significant speed gain.

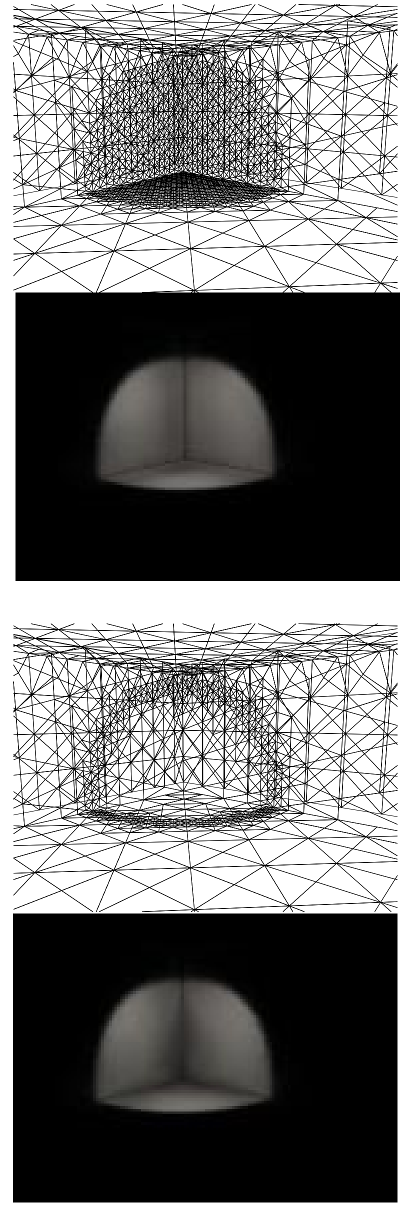

a

b

Figure 7.3.2.2: (a) Overmeshing due to a low **Threshold** parameter. (b) Meshing corrected with an increased **Threshold** parameter.

The image appears over or underexposed

Radiosity models light distribution in an environment in a physically accurate manner. Given that the intensity of light sources is specified in accurate units, such as Watts or Lumens, intensities recorded at mesh polygons represent a physical intensity unit. While the human eye is capable of adjusting to a wide range of light intensities, a display medium such as the computer screen is not. That is, while the intensities at the mesh polygons are physically accurate, when mapped directly on the screen, the inability of the screen to display very bright intensities or very dim intensities in a visible manner leads to over or underexposure of the image

An image may also contain areas with both very bright or very dim areas. If the intensities are mapped on the screen in a linear fashion, the bright areas may appear overexposed, while the dim areas may appear underexposed. This can be corrected by mapping mesh intensities into the intensity range which can be handled by the computer screen. Selecting the **Image Exposure...** button in the image options dialog invokes the **Image Exposure Options** dialog, where the intensity range and mapping of physical intensities to screen intensities can be defined. An example of an image with a corrected exposure is shown in Figure 7.3.2.3.

a

b

Figure 7.3.2.3: (a) Uncorrected and
(b) corrected exposure of a radiosity based rendering.

The radiosity based image does not show any effects from bump mapping

Bump mapping is an effect which is created by the intentional disturbance of the surface normal direction during the shading calculation of the simple illumination model. The accurate representation of diffuse reflection during the radiosity solution cannot take into account such surface disturbances. As a result bump mapping effects in a rendering are only created by specular reflections, mirrored reflections and refracted transparencies, which are still calculated on a per pixel basis during the rendering.

The progress bar does not advance during the execution of the radiosity solution

The execution of the radiosity solution is not a linear process. That is, the first 60% of the solution may be achieved very quickly. The amount of time it will take to advance from 60 to 70% may be much longer and even longer to get from 70 to 80%. As the solution approaches 100%, which it never will, more and more time will be spend to gain less and less accuracy. While it appears, that the progress bar does not advance, the computational effort is spent on calculating a minimal gain of the solution, which cannot be represented in the progress bar. This is especially the case, when the mesh density is high and the solution has advanced beyond an accuracy of 80%. In many cases, the visual gain of a solution which is more than 85% accurate is small, and the radiosity process may be interrupted manually by clicking in the **Stop** button of the progress bar.

The rendering based on a radiosity solution is all black

There are a number of reason why this may occur. First, it is possible that the light sources are simply not strong enough to illuminate the scene. Recall that in the real world the intensity of light falls off with the square of the distance the light travels. That is, a point which is twice as far away from the light source than another point only receives a quarter of the intensity of the closer point. For example, if a candle is placed inside an auditorium, the walls of the space are simply too far away to receive any visible light energy. It is important, to adjust the intensity of the light with respect to the size of the space which is illuminated. Unlike the simple illumination model, which allows for a constant, linear or square intensity falloff, the radiosity process always assumes square falloff of light intensities.

Another reason for non illuminated scenes may be, that light, which shines from outside into a space through openings in the walls, is obscured by objects in the openings. If a window is modeled with an object representing the glass surface, the object must have the shadow casting attribute turned off, even if the surface style used by the object has transparencies. Unlike the simple illumination model, radiosity currently assumes all shadows to be opaque.

7.4 Radiosity tutorial

The tutorial in this section will cover how the image in Figure 7.4.0.1 was rendered, using a radiosity solution. It is an indoor scene sometime in the early evening.

You will find the project file for this tutorial on the product CD-ROM. It is in the “Tutorial” folder and is called “nolite.fmz”. Since this tutorial is about radiosity solutions, we recommend that you use our model, rather than re-doing it. At the same time, we recommend that you set the lights yourself, rather than using the file on the CD-ROM that already has them. The six lights you will be setting are shown in Figure 7.4.0.2.

This tutorial consists of seven subsections, one for each of the lights that we will be setting up, and one for the final manipulation of our light parameters. Each subsection displays the light to be set, the dialogs where the settings are made, and, at the end, the rendering that results from the inclusion of the light. These renderings are cumulative; that is, each time a new light is added the previous lights are also included. Once all the lights have been included, we experiment with refinements of the initial rendering.

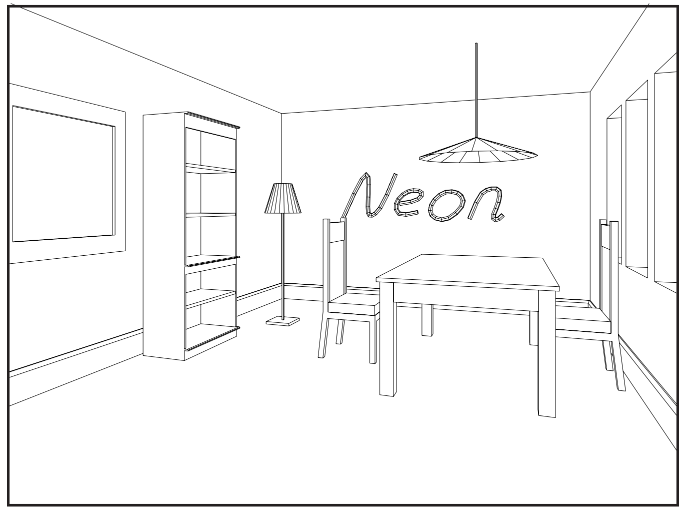

a

b

Figure 7.4.0.1: (a) A scene in wire frame and (b) a radiosity based rendering.

Since computing a highly accurate radiosity solution tends to be a fairly time consuming process, it is best to start with radiosity parameters set very low. As the radiosity results become more finalized, increasingly accurate solutions may be used. Note that, as a general rule, the lower the number of faces in your scene the faster you will arrive at a satisfactory radiosity based rendering. It is also helpful to work with only one light active at a time, as you set up a scene's lighting.

Throughout this tutorial, we will provide approximate times you may expect to wait in order to generate radiosity solutions for each step. These times were based on the time it took to generate solutions on a Power Mac 8500/120.

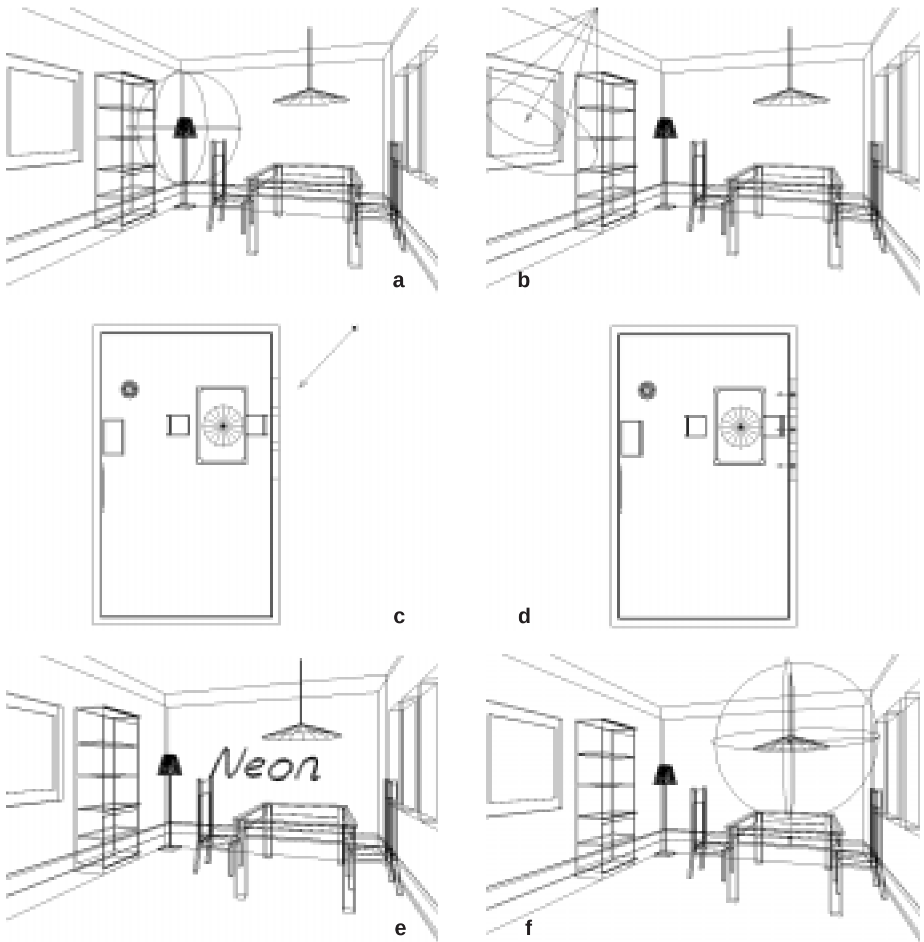

Figure 7.4.0.2: The lights in the scene: (a) point, (b) cone, (c) sun, (d) and (e) area, and (f) custom.

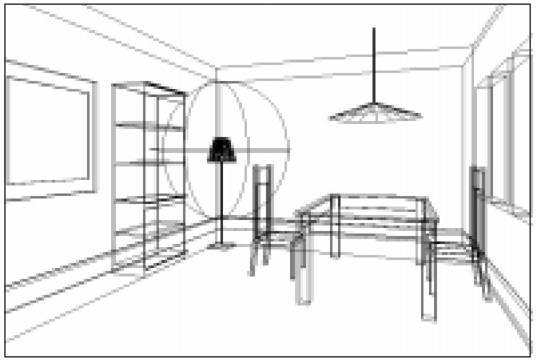

Figure 7.4.1.1: The point light to be defined.

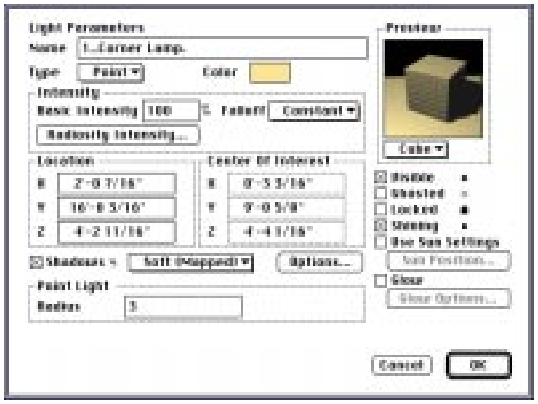

Figure 7.4.1.2: Creating the point light.

Conversion chart from watts to lumens

Watts	25	40	60	75	100	150	200
Lumens	230	430	730	960	1380	2220	3150

Figure 7.4.1.3: The watts-lumens equivalencies for inscandescent lights.

7.4.1 Point Light

The first light we will create is a point light for the corner lamp (Figure 7.4.1.1). We use a point light in this case because we can achieve an even distribution of light from a single point source. We let the lamp shade control how the shadows will be cast from this light. Before you do anything confirm that the shinning attribute of the sun light that already appears in the Lights palette by default is off.

• Create a new light by clicking under the last entry of the **Lights** palette. Double click on the new light name to invoke the **Light Parameters** dialog.
• In the **Light Parameters** dialog, select **Point** from the **Type** pop up menu.
• Set **Radius** to 3'-0".
• Set the **Color** to yellow and make sure that **Shadows** is on (Figure 7.4.1.2).

It should be noted that the **Basic Intensity** value in the **Light Parameters** dialog has no effect when generating a radiosity solution. This value does, however, determine the intensity of specular highlights that are generated through the **RenderZone*** rendering of a radiosity solution. If a radiosity solution is not being calculated before a **RenderZone*** rendering, this value will determine the overall intensity of the light, not just the specular intensity.

It should also be mentioned that the type of shadow used is irrelevant for the radiosity solution. The type of shadow, soft or hard, will only affect the specular highlights that may be generated during the **RenderZone*** rendering portion of a radiosity based image.

Recall that we can use either a **Radiometric** (**Watts**) or a **Photometric** (**Lumens**) method for setting the intensity of a point light. While later in this tutorial we mostly use **Watts**, in this section we shall use **Lumens**. We only do this to offer an example of using these units of light measurement. We have also provided a small conversion chart (Figure 7.4.1.3) that shows the appropriate values in lumens for several wattages of common incandescent light bulbs. For a more complete conversion table see section 7.2.6.

• In the **Lights Parameters** dialog, click on the **Radiosity Intensity...** button to invoke the **Radiosity Intensity: Cone And Point** dialog, as shown in Figure 7.4.1.4.

• Select **Photometric** and set the **Brightness** to 730 Lumens. This setting simulates the light intensity of a typical 60 watt light bulb.

Click **OK** twice to return to the model window and move this light into place inside of the shade.

A quick radiosity solution may be generated at this time with just the point light active. Remember as we progress through this tutorial, successive lights will be added to each new solution.

• Set display mode to **Shaded Render***.
• From the **Display** menu, select **Radiosity Options...** to open the respective dialog.
• In the **Radiosity Options** dialog set **Preset Parameters** to **3** to generate a quick yet accurate solution (Figure 7.4.1.5).
• Click **OK** then **Generate Radiosity Solution*** to start the radiosity solution process. The results should be roughly as in Figure 7.4.1.6.

• Select **Exit Radiosity** (**Display** menu).

Note that **Apply Reflected Light As Ambient** in the **Radiosity Options** dialog is on by default and we left it on. Thus the portion of light that has not yet been calculated by the radiosity solution is applied to the rendering as ambient light. We keep this option on throughout this tutorial.

The time for this “quick” solution should be no more than a minute (on an 8500/120 Power Macintosh). All the quick renderings that follow in this tutorial were also under 5 minutes. Only at the very end we allow longer rendering times.

If you wish to make changes to the light’s intensity or other parameters, make the changes, then regenerate a new radiosity solution. This can be done repeatedly until the results are satisfactory.

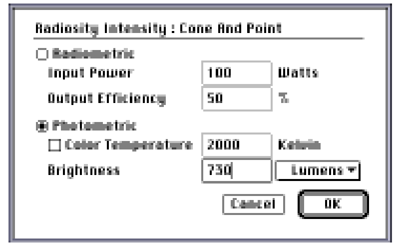

Figure 7.4.1.4: Setting the **Radiosity Intensity** parameters.

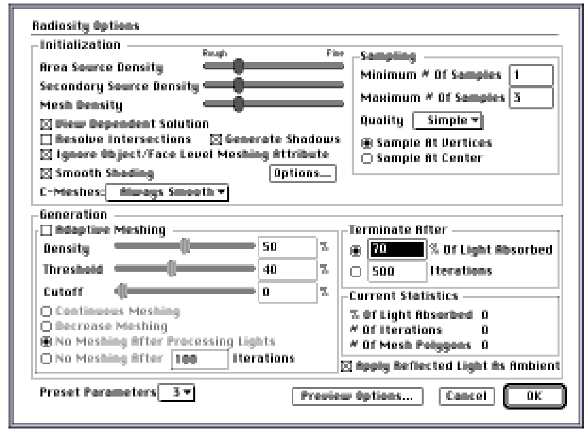

Figure 7.4.1.5: Setting the **Radiosity Options**.

Figure 7.4.1.6: A radiosity based rendering of the scene with the point light active.

Figure 7.4.2.1: The cone light.

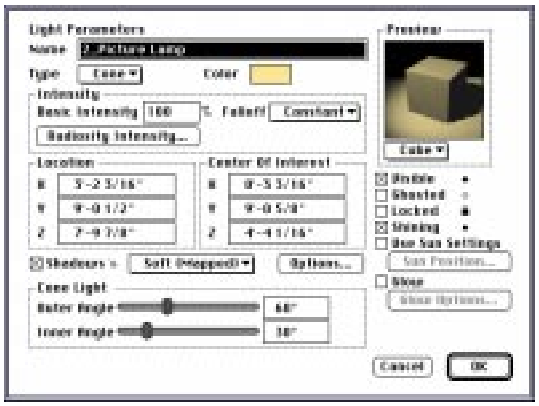

Figure 7.4.2.2: Creating the cone light.

Figure 7.4.2.3: A radiosity based rendering of the scene after adding the cone light.

7.4.2 Cone Light

Before proceeding with setting the next light, we need to do a couple of things. Execute **Wire Frame*** from the **Display** menu and turn off the visibility attribute of the previous light. This will let you see better what you will be doing next. Keep on doing these adjustments at the beginning of all the sections that follow.

Next we will create a cone light for the picture on the wall (Figure 7.4.2.1). We use this type in order to have directional control of the light. Since we do not want to wash out the picture, we will use fairly low intensities with soft shadow boundaries.

• Create a new light and set the **Type** to **Cone**.
• Set **Color** to a light yellow and chack that **Shadows** is on.
• Set the **Inner Angle** to roughly half the **Outer Angle** to generate soft shadow boundaries (Figure 7.4.2.2).
• Click on the **Radiosity Intensity...** button to invoke the **Radiosity Intensity: Cone And Point** dialog.
• Select **Radiometric** and set the **Input Power** to 20 watts and the **Output Efficiency** to 50%.
• Click **OK** twice to exit the dialog.
• Move the cone light so that it shines down on the picture from a point near the ceiling.
• Set view mode to **Shaded Render***.
• **Generate Radiosity Solution*** to check the results. Use the same settings in the **Radiosity Options** dialog (**Preset Parameter** of 3) that were used for the point light.

The results should be roughly as in Figure 7.4.2.3.

• Select **Exit Radiosity**.

Make whatever changes to the light parameters you wish, until you are satisfied with the results.

7.4.3 Distant Light

This light will simulate light from the outdoor environment, shining through the windows and onto the floor. This will be the already existing sun light, whose shining attribute we now need to turn back on. After you turn it on, set it in such a way that it is directed towards the windows (Figure 7.4.3.1). If the distant light was not directed towards the window, the room would not be lit with this light at all. Note that we shall not use the **Atmospheric Light** option for this light (see below).

- Double click on the distant light name in the Lights palette to invoke the **Light Parameters** dialog.
- For this light to cast shadows on the floor, turn on **Shadows** and set the **Color** to white (Figure 7.4.3.2).
- Click on the **Radiosity Intensity** button to invoke the **Radiosity Intensity: Distant** dialog.
- Set **Output Power** to .25 watts per square foot. Note that typically this value would be about 2 watts per square foot. Because we are assuming an early evening sun, we use relatively low intensities.
- Click **OK** twice to return to the model window.
- Set display mode to **Shaded Render***.
- **Generate Radiosity Solution*** and check the results. They should be roughly as shown in Figure 7.4.3.3.
- Select **Exit Radiosity** to exit.

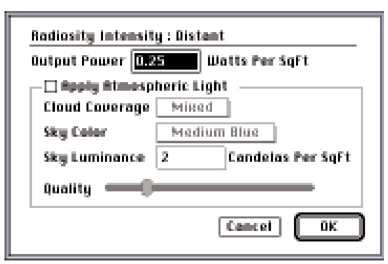

Figure 7.4.3.4: The **Apply Atmospheric Light** option in the **Radiosity Intensity: Distant** dialog.

Note the light shining on the floor and recall the option **Apply Atmospheric Light**, available in the **Radiosity Intensity: Distant** dialog (Figure 7.4.3.4). Had we used this option, an overall light would be added, but there would have been no lighted area on the floor.

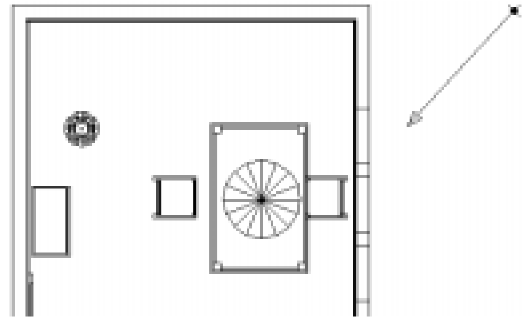

Figure 7.4.3.1: The sun light.

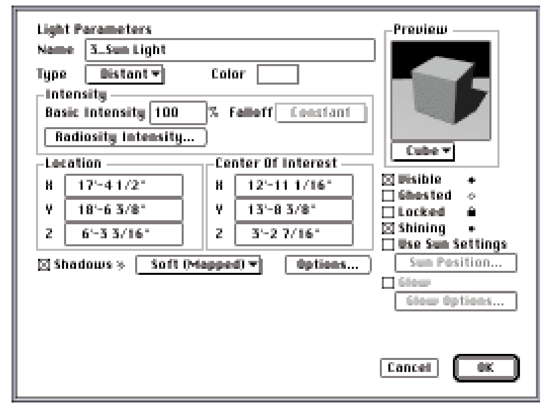

Figure 7.4.3.2: Setting the parameters of the sun light.

Figure 7.4.3.3: Radiosity rendering of the scene after adding the sun light.

7.4.4 Area Lights

Area lights work by turning an object into a light. In this tutorial example, the surface objects representing the window panes will be converted into area lights. These lights will work in conjunction with the distant light and will complete the effect of light coming through the window. They will simulate the diffuse light that indirectly enters through the window, which is light reflected from clouds, buildings, the ground, etc. This light will illuminate areas of the room that are in close proximity to the windows.

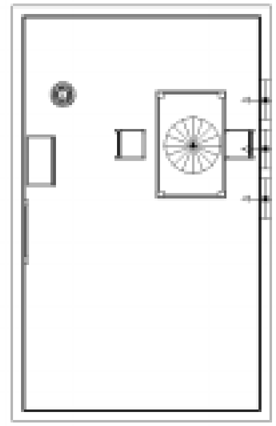

Figure 7.4.4.1: The window panes transformed into area lights.

The surface style of the window pane needs to be set before the object is converted to an area light. This is because, once an object is converted to an area light, its attributes can no longer be modified.

If the geometry or the object attributes of an area light need to be modified, the original object has to be recovered from the area light. This is done by changing the area light to another light type. Once the changes have been made, the object can again be transformed to an area light.

Also, the window panes must have their shadow casting attribute turned off, because radiosity assumes all shadows to be opaque. If the window panes cast shadows, they will block out the light generated from the distant light source that is located outside of the room.

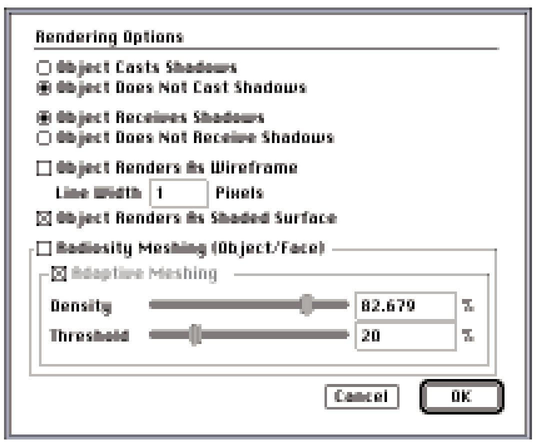

Figure 7.4.4.2: Selecting settings in the **Rendering Options** dialog.

1. After the surface style for the windows is set, we also need to set the windows' rendering attributes, so that they do not cast shadows and block the distant light.

- Using the Pick tool, prepick the window pane objects.
- Double click on the Render Attributes tool to invoke its dialog.
- From the **Rendering Options** dialog, select **Object Does Not Cast Shadows** (Figure 7.4.4.2).
- Click **OK** to exit the dialog.
- With the Render Attributes tool active, click anywhere in the graphics window to apply the rendering attributes to the preselected objects (pane windows).

2. The directions of an area light coincide with the directions of the face normals of the object that will be made into an area light. Consequently, we can display the normals of an object in order to check the direction towards which it will be shinning, after it is made an area light.

• Turn on **Show Face Normals** in the **Wireframe Options** dialog. If they do not have the desirable direction, to reverse them, double click on the Topo Attributes tool to invoke its dialog, and in it select **Reverse Direction**.
• After exiting the dialog, with the Topo Attributes tool active, click on the window pane whose normal you wish to reverse. The normals of the window panes should point towards the inside of the room, as shown in Figure 7.4.4.1.

3. Since these lights are going to be used for very subtle lighting effects around the window openings, very low intensities will be used.

• With the Pick tool, preselect a window pane.
• Click under the last entry of the Lights palette to create a new light. Invoke the **Light Parameters** dialog.
• In the **Light Parameters** dialog, select **Area** from the **Type** pop up menu.
• Select **Render Light Object**, which causes the window pane to be rendered with the surface style previously set (Figure 7.4.4.3).
• Click on the **Radiosity Intensity...** button to invoke the **Radiosity Intensity: Area** dialog (Figure 7.4.4.4). In it select **Radiometric** and set the **Output Power** to 6 watts. Click **OK** to exit the dialog.
• Exit the **Light Parameters** dialog.

Repeat these steps for each of the window panes.

• Select **Shaded Render***(**Display** menu).
• **Generate Radiosity Solution***. The result should be as in Figure 7.4.4.5).
• Select **Exit Radiosity**.

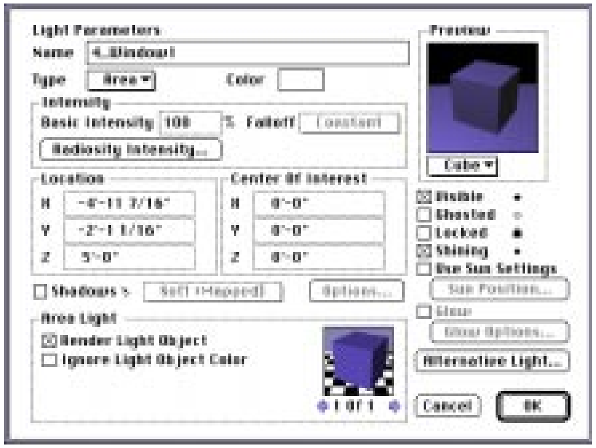

Figure 7.4.4.3: Setting the parameters of the window area ligths.

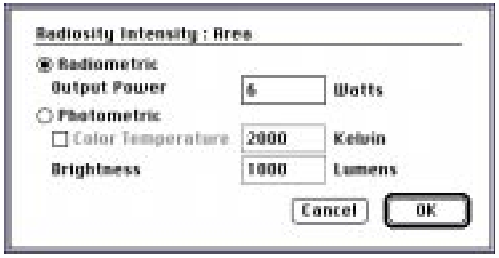

Figure 7.4.4.4: Setting the **Radiosity Intensity** options for the window area lights.

Figure 7.4.4.5: A radiosity based rendering of the scene after adding the window area lights.

7.4.5 One more area light: Neon

The next light will be the neon sign, which is again an area light. Obviously the “Neon” object, which is a group of sweeps, needs to be constructed first. Then it can be transformed into a light.

- Preselect the neon object using the Pick tool.
- Click under the last entry of the **Lights** palette to create a new light.
- In the **Light Parameters** dialog, select **Area** from the **Type** pop up menu.
- Set the **Color** to a bright red and turn on **Render Geometry** so that the sign itself will render (Figure 7.4.5.2).

Note that the color of an area light is a mixture of the object’s surface style color, before converting it to an area light, and the light’s color. By modifying either color, you can control the color emitted from the area light. For example, a red area light results from a white surface style set for the object and a red color set for the light. You can produce the same result by setting the surface style color to red and the light color to white.

- Invoke the **Radiosity Intensity: Area** dialog and in it select **Radiometric** and set **Output Power** to 8t (Figure 7.4.5.3). Exit the dialog.

At this point you can again generate a new radiosity solution with all the current lights active. The results should be as shown in Figure 7.4.5.3.

As you have noticed, this rendering took a bit longer due to the additional lights it now has to process. This time should still be under 5 minutes.

Figure 7.4.5.1: The Neon area light.

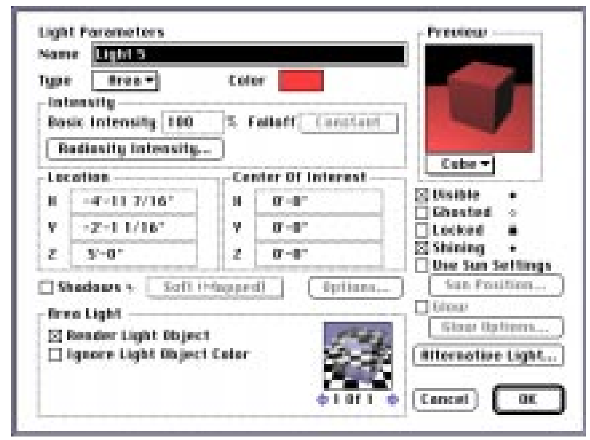

Figure 7.4.5.2: Setting the parameters of the Neon area light.

Figure 7.4.5.3: A radiosity based rendering of the scene after adding the Neon area light.

7.4.6 Custom Light

For the light over the table we shall use a custom light, which allows us to control exactly how the light falls off from its source.

- Create a new light and invoke the **Light Parameters** dialog.
- In the dialog, select **Custom.**
- Enter 0,0,4’ for **Location** and 0,0,0 for **Center of Interest**. These values will set the center of interest directly below the origin of the light.
- Set **Color** to yellow and confirm that **Shadows** is on (Figure 7.4.6.2).

- Invoke the **Radiosity Intensity: Custom** dialog and in it select **Radiometric**, set the **Input Power** to 40 watts, and the **Output Efficiency** to 40%. These values roughly correspond to a conventional incandescent light.
- Set **Symmetry** to **Axial** and **Hemisphere** to **Bottom**. We use these settings because the light is symmetrical about its source and the lamp shade dictates that no light will shine in an upwards direction. When you select **Bottom**, a warning is posted. Click **OK** and proceed.
- Set the light curve as shown in Figure 7.4.6.3. Recall that you can insert points by clicking on the curve while pressing the ***option*** (Macintosh) or ***shift+ctrl*** (Windows) keys. The shown shape roughly corresponds to the light distribution curve of an actual light bulb.
- Click on **Alternative Light...** and, in the dialog that is invoked, set **Radius** to 3’-0”. This sets the radius of the sphere that represents the custom light in wire frame (in addition to defining the parameters of the light that will be used to approximate the custom light when rendered outside radiosity).
- Exit the dialog, pick the custom light, and move it into position under the lamp shade, above the table.

Figure 7.4.6.1: The table custom light.

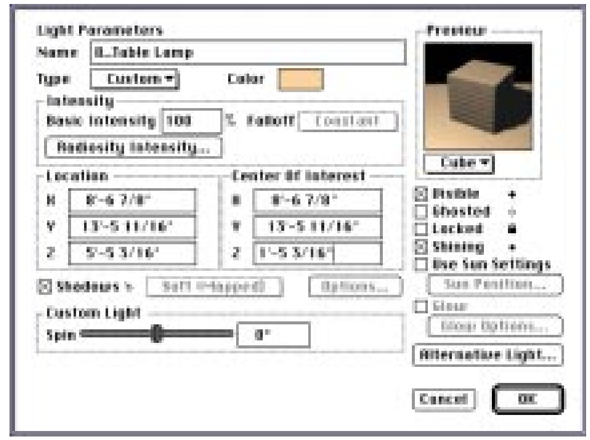

Figure 7.4.6.2: Setting the parameters of the table custom light.

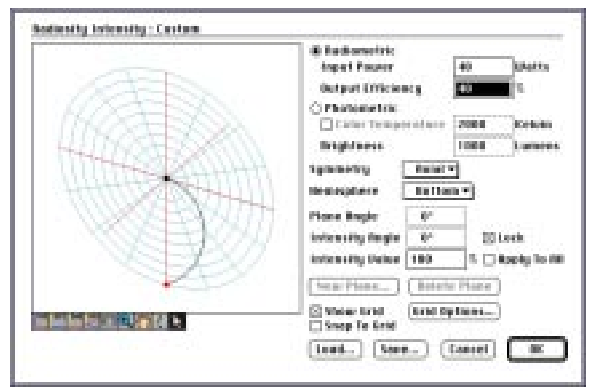

Figure 7.4.6.3:
Setting the **Radiosity Intensity: Custom** parameters for the table light.

At this point you may generate a solution with all the lights active to see if the results are appropriate.

- Set view mode to **Shaded Render***.
- From the **Display** menu, select **Radiosity Options...** to invoke the **Radiosity Options** dialog.
- Set **Preset Parameters** to **3** to generate a fairly quick and accurate solution.
- Exit the dialog and **Generate Radiosity Solution***.

The results should be roughly as shown in Figure 7.4.6.4. Time for this solution was approximately 5 minutes.

- Select **Exit Radiosity**.

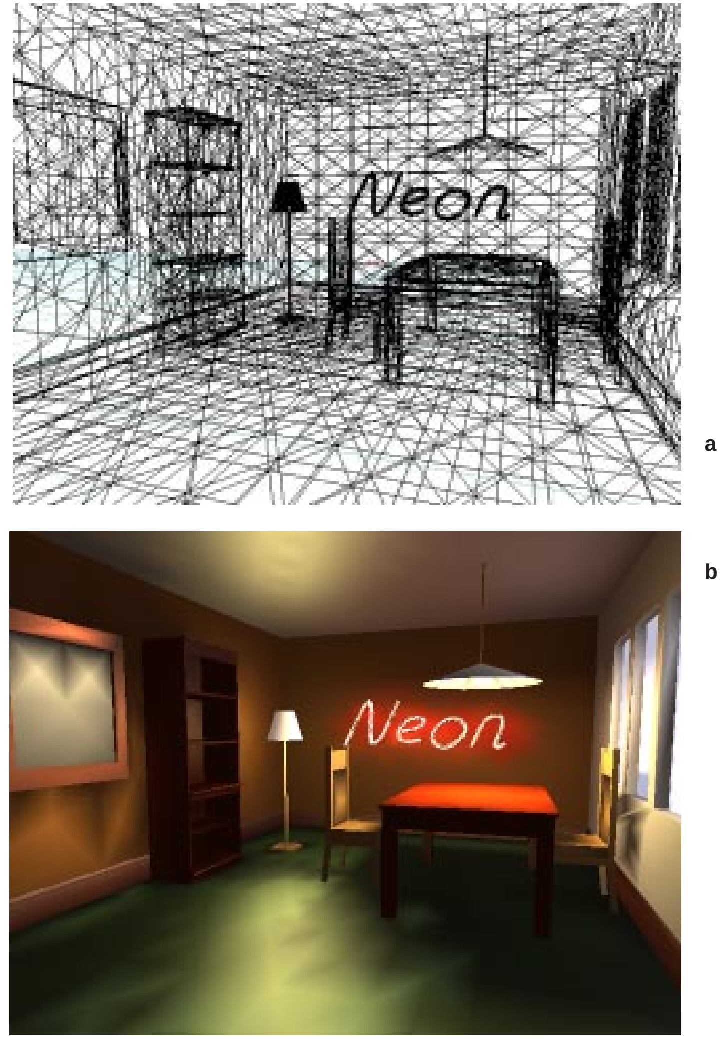

Figure 7.4.6.4: (a) A radiosity mesh of the scene and (b) a rendering based on it, after adding the table custom light.

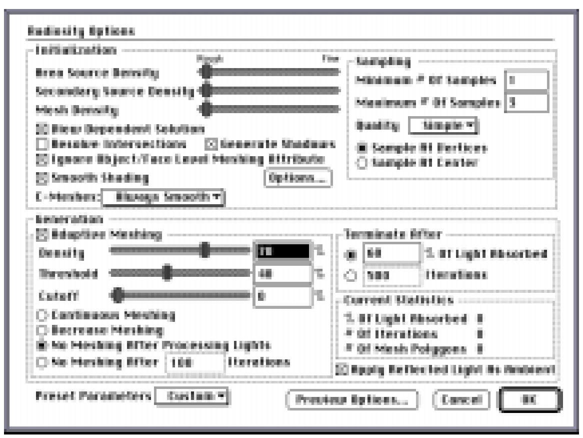

Figure 7.4.6.5: Adjusting the radiosity parameters.

All the radiosity renderings produced up to this point were based on low resolutions. Before completing this section we shall produce a higher density radiosity solution. Assuming that the all the light intensities are still set properly, as suggested in the previous sections, to generate a better solution we shall simply select parameters that create a denser mesh. The new rendering will offer us a better understanding of where additional refinements may be needed. These refinements are applied in the next section.

After setting the radiosity parameters as shown in the dialog in Figure 7.4.6.5, repeat the previous process. The solution is again generated fairly quickly but is of a better quality, as shown in Figure 7.4.6.6. This is due to the denser mesh along the shadow edges that is now produced. However, even this is not yet perfect. We shall further refine our radiosity based rendering in the next section.

a

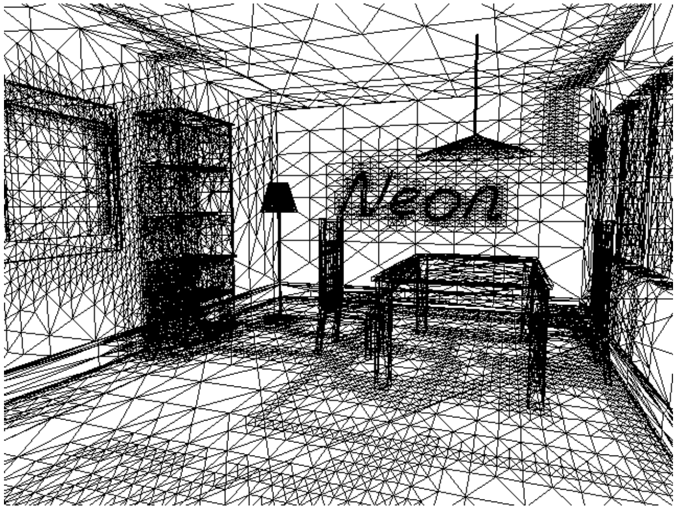

b

Figure 7.4.6.6: (a) A denser radiosity mesh and (b) a rendering based on it, with all the lights active.

7.4.7 Refining the radiosity solution

Looking at the previous radiosity rendering (Figure 7.4.6.6), we can see that the floor and the walls have jagged shadow edges. Next, we shall refine the solution to generate better shadows. We shall optimize the radiosity mesh so that areas with higher contrast will have a denser mesh than areas with more subtle changes in light intensity. We shall do this using two techniques: (i) pre-meshing some of the faces in our scene, and (ii) applying radiosity parameters at the object level, rather than applying them globally only.

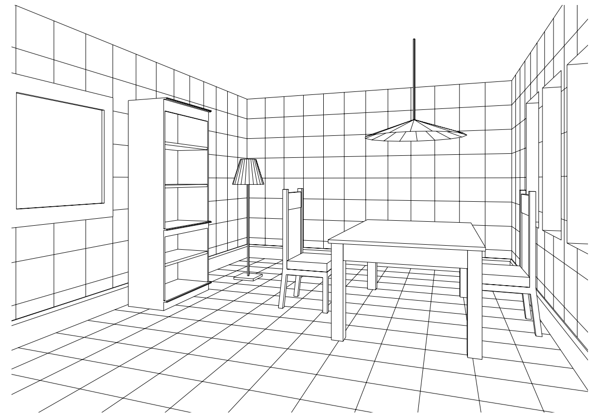

Figure 7.4.7.1:
Pre-meshing the faces of the floor and walls.

1. The purpose of pre-meshing the faces of the walls and floor is to provide a denser and more even mesh for the radiosity solution to subdivide. Large face areas may subdivide when generating a radiosity solution in undesirable and unpredictable ways, unless we control the size of the faces that shadows will fall on. By pre-meshing those areas that will have shadows cast upon them, we are able to direct how shadow edges will subdivide.

- Double click on the Mesh tool to invoke its dialog and in it set mesh intervals to 1'.
- Set topological level to Face and, with the Mesh tool active, select the faces of the walls and the floor. Note that the orientation of the mesh is significant and it should follow the direction of the shadows. In this case, the majority of the shadows follow the direction of the room. The mesh should be as in Figure 7.4.7.1.

2. Recall that the parameters that are set in the **Radiosity Options** dialog have a global effect. In other words, those settings affect all the objects in the scene. You can also apply these parameters to objects individually, which offers significant optimization possibilities.

Areas with well defined shadow edges will be meshed the same as areas with softer shadow boundaries. A major part of the radiosity optimization process involves setting the areas with well defined shadow edges to a higher density mesh, while using the global settings in areas with more subtle changes in light, to generate a lower density mesh. This way the solution takes less time and the results are the same as if a denser mesh were generated globally.

Looking at our previous rendering (Figure 7.4.6.6) we can see that the floor has the coarsest shadow edges. Consequently, we shall set these faces to a higher mesh density and a lower threshold while meshing the rest of the scene at a lower density and higher threshold. Also, underneath the picture on the wall there are some rough shadow edges. We can select the faces that are directly under the picture and set them with the same density and threshold as the floor.

• Preselect all the floor faces and the faces directly beneath the picture.

• Double click on the Render Attributes tool to invoke the **Rendering Options** dialog. In it select **Radiosity Meshing**, **Adaptive Meshing**, set **Density** to 80%, and set **Threshold** to 10%, as shown in Figure 7.4.7.2. Click **OK**.

• With the Render Attributes tool active, click anywhere in the window to apply the adaptive meshing attributes to the preselected faces. These settings will override those in the **Radiosity Options** dialog.

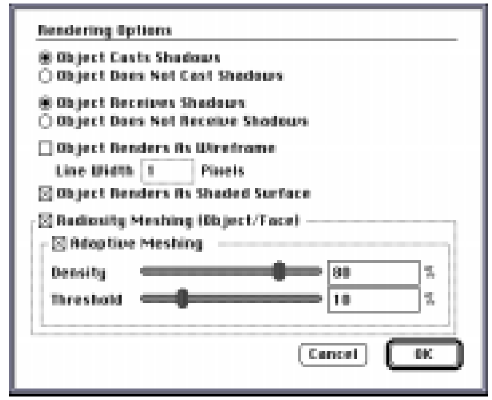

Figure 7.4.7.2: Setting the **Rendering Options** parameters.

At this point we are ready to try a more accurate solution. We set the parameters in the **Radiosity Options** to lower settings than the settings we just used on the individual faces and still arrive at a satisfactory radiosity solution.

• Click on **Radiosity Options...** to invoke its dialog and from the **Preset Parameters** pop up menu select 6.

• Since we shall only need one view of the scene, you may want to turn on **View Dependent Solution** to save some time.

• Increase the **Adaptive Meshing Density** to 75% and decrease the **Threshold** to 20%.

• Set to **Terminate After** 75% **Of Light Absorbed.**

• Set display mode to **Shaded Render***.

• **Generate Radiosity Solution***.

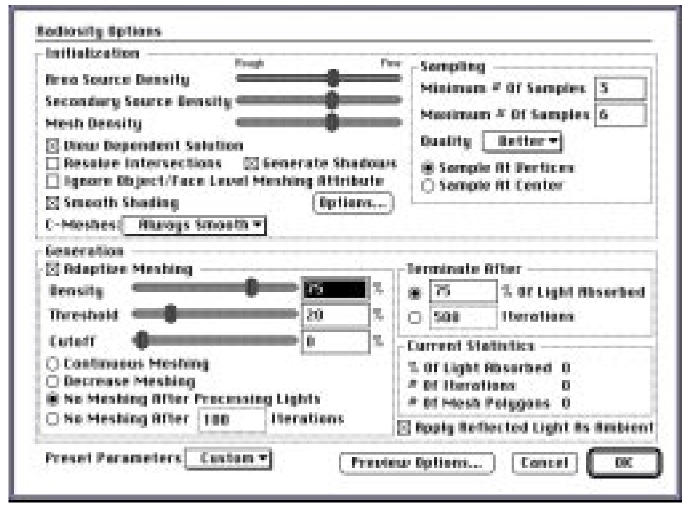

Figure 7.4.7.3:
Setting the global radiosity parameters.

When a satisfactory radiosity solution has been generated, do a **RenderZone*** rendering at **Full Raytrace,** without exiting the solution. This should produce an image roughly as shown in Figure 7.4.7.4. The time to generate this solution was approximately 4 hours.

Figure 7.4.7.4: A refined radiosity based rendering.