materials. One way to incorporate complex lighting effects, such as volumetric shadows, is to precompute a shadow volume for storing the amount of light arriving at each sample after being attenuated by the intervening material. During rendering, the interpolated values from this volumetric shadow map are multiplied by colors from the transfer function. But in addition to using extra memory, volumetric shadow maps result in visual artifacts such as blurry shadows and dark images.

A better alternative is to use a pixel buffer to accumulate the amount of light attenuated from the light's point of view (Kniss et al. 2003). To do this efficiently, the slicing axis is set halfway between the view and the light directions, as shown in Figure 39-9a. This allows the same slice to be rendered from both the eye and the light points of view. The amount of light arriving at a particular slice is equal to 1 minus the accumulated opacity of the previously rendered slices. Each slice is first rendered from the eye's point of view, using the results of the previous pass rendered from the light's point of view, which are used to modulate the brightness of samples in the current slice. The same slice is then rendered from the light's point of view to calculate the intensity of the light arriving at the next slice. Algorithm 39-3 uses two buffers: one for the eye and one for the light. Figure 39-10 shows the setup described in Algorithm 39-3.

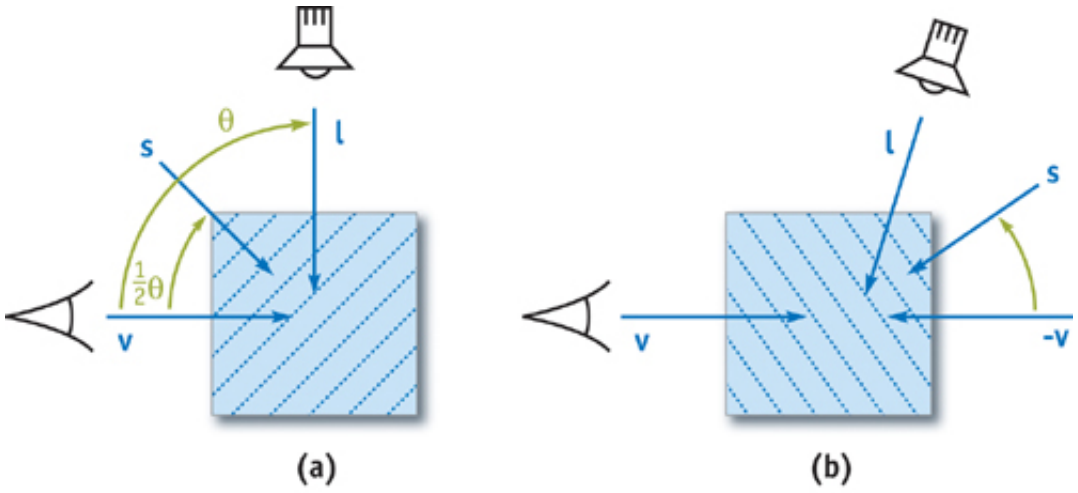

Figure 39-9 Half-Angle Slicing for Incremental Lighting Computations