not drawing empty regions of the volume (Li et al. 2003, Krüger and Westermann 2003). Empty-space skipping can efficiently balance the available geometry and fragment processing bandwidth.

Fragment Program Limitations

Volume rendering performance is largely influenced by the complexity of the fragment shader. Precomputed lookup tables may be faster than fragment programs with many complex instructions. Dependent texture reads, however, can result in pipeline stalls, significantly reducing rendering speed. Achieving peak performance requires finding the correct balance of fragment operations and texture reads, which can be a challenging profiling task.

Texture Memory Limitations

The trilinear interpolation (quadrilinear when using mipmaps) used in volume rendering requires at least eight texture lookups, making it more expensive than the bilinear (trilinear with mipmaps) interpolation used in standard 2D texture mapping. In addition, when using large 3D textures, the texture caches may not be as efficient at hiding the latency of texture memory access as they are when using 2D textures. When the speed of rendering is critical, smaller textures, texture compression, and lower precision types can reduce the pressure on the texture memory subsystem. Efficient compression schemes have recently emerged that achieve high texture compression ratios without affecting the rendering performance (Schneider and Westermann 2003). Finally, the arithmetic and memory systems in modern GPUs operate on all values in an RGBA tuple simultaneously. Packing data into RGBA tuples increases performance by lessening the bandwidth requirements.

39.7 Summary

Volume rendering is an important graphics and visualization technique. A volume renderer can be used for displaying not only surfaces of a model but also the intricate detail contained within. The first half of this chapter presented a typical implementation of a texture-based volume renderer with view-aligned proxy geometry. In Section 39.5, two advanced techniques built upon the basic implementation were described. The presented techniques improve the quality of images by adding volumetric shadows, translucency effects, and random detail to the standard rendering model.

Volume rendering has been around for over a decade and is still an active area of graphics