Amplitude Interference Modeling

by Patrick Lacz
based on “Wavelength Dependent Reflectance Functions”
by Jay Gondek, Gary Meyer, and Jonathan Newman
of the University of Oregon

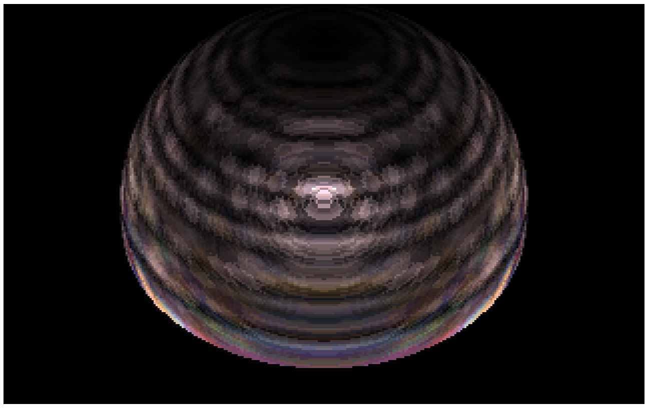

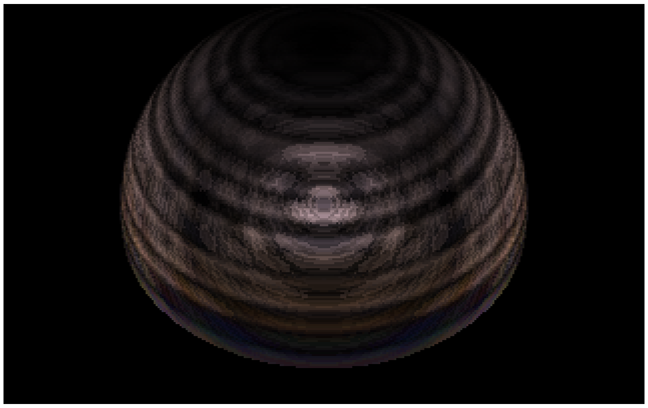

Amplitude interference is an important part of appearance modeling for a remarkably large number of real-world items. The effect is most commonly associated with gasoline, bubbles, and tinted glass, however it is also an important aspect of many paints. The authors of the paper that I have based this work upon also used it as a principle motivation for developing a virtual goniospectrometer and a supporting BRDF representation.

At the heart of my implementation and theirs is the virtual goniospectrometer. In real life, a goniospectrometer takes a small sample surface and determines lighting characteristics based on a large number of simulated lighting conditions. Inside the computer, the process is the same: take a small sample of a surface and record the reflectance in a large number of lighting conditions.

The surface sample is usually developed independently of the final scene and contains a much more detailed representation than a more macroscopic renderer would desire. The microgeometry that many BRDF models approximate are made explicit in the surface sample. Paint, for instance can be modeled as a large number of pigmented spheres possibly suspended in some other medium. Depending on how dry the paint is, the spheres may or may not change the roughness of the surface.

As mentioned, one of the principle motivations for simulating the microgeometry is modeling amplitude-splitting interference. In this form of interference, geometry details within the wavelength of the visible spectrum can cause cancellation due to phase differences along the wave. Only by ray-tracing within a microgeometry scene can these details be handled arbitrarily.

The major hurdle in any implementation of this form is the data structures. The principle contribution of the paper is a quad-tree based subdivided BRDF representation which is used to accumulate the multitude of samples into a queriable data structure.

When working to implement their structures, I became concerned that the structure was not fully specified, may not work as described and at the very least was overly complex. I set off to make a simpler structure based on a kd-tree.

Initially, I reasoned that superposition should be handled per-BRDF query, so that fewer sample points would be required. The changes to the kd-tree to support proper scaling of the results became prohibitive, and the required sampling density was by no means small. Generating over a million samples still produced poor results where interference effects were not visible.

I had been using a 6-dimensional kd-tree representing the BRDF query as two unit vectors. This allowed me to ignore the non-linear spherical coordinates of the 4-d BRDF space. I decided to switch representations. First, I generalized the kd tree to use a separate distance function for each dimension. The distance was measured in the sum of the arc-lengths distance for the incident and exitant illumination. In order to find the arc-length distance to a chosen pivot azimuth, the arc-length to a point reflected about that azimuth is halved.

I also added a reduction phase to the data-structure generation. No longer were the samples taken during microstructure simulation the input data for BRDF queries. I descritized the space into buckets and accumulated samples into each. Superposition was applied to samples mapping to the same bucket. Each bucket then was entered as an entry into the new 4-d spherically parameterized kd-tree. During BRDF evaluation, the nearest buckets to the queried point are found and superimposed to produce the final result.

The new structure allowed me to generate a better estimate for the number of rays cast from the incident direction, which in turn allowed for a more accurate reflectance estimate. Superposition was carried out on two levels: when creating and combining the buckets. As the sample images show, this still allows highly varying cancellation effects to occur nearby.

The major disappointment of the new "simplified" implementation is that it still does not scale nicely based on more through sampling of the microgeometry. The left image at the top of this paper is formed using approximately 45000 sampled locations in the BRDF and assuming relative isotropy. The microgeometry is modeling a thin film with an average thickness of 600 nm. The image on the right is made from roughly 10 times as many samples and shows significant darkening effects, which I believe to be caused by an excess of cancellation that I have not yet identified.