Andrew Glassner's Notebook

http://www.research.microsoft.com/research/graphics/glassner/

Solar Halos and Sun Dogs

Andrew Glassner

Microsoft

One of the things I love about computer graphics is the sheer variety of topics we work with. When I started studying computer graphics, I expected to learn some physics, mathematics, and computation, but gradually I discovered that our field embraces physiology, art technique, perception, and more. I continue to enjoy graphics because I love the steady flow of new ideas drawn from the sciences and the arts.

In this column, my goal is to bring up topics I think you will find interesting and discuss how they make use of, and sometimes extend, traditional computer graphics tools. I will try to discuss topics that the typical *CG&A* reader will find intellectually stimulating. I will often skip over the exhaustive detail necessary in formal technical papers, leaving it to you to fill in the blanks (but I'll try to make sure the blanks aren't too big). In the same spirit, I won't present many detailed algorithms or general graphics tools; Jim Blinn does a splendid job of that in his excellent column.

I'll try to provide a lot of information, but I'll leave many issues for you to think about (sometimes I'll point them out along the way). I'll also draw on whatever graphics tools are appropriate to the job at hand. If I seem to take for granted something from computer graphics that's unfamiliar to you, it can probably be found in the standard texts. I'll also provide a bibliography with each column for further reading on that column's subject.

My inaugural topic is something with which we are all familiar: the beautiful displays of light in the air around us. More specifically, I will talk about the elegant mechanisms that bring about the striking atmospheric effects associated with solar halos.

Solar halos

If you walk outside on a sunny winter day when there is moisture in the air, you might notice a bright circle surrounding the sun, as in Figure 1. This is a solar halo: it has a sharp inner edge located about 22 degrees away from the center of the sun, then it fades out. Sometimes only fragments of the halo are visible, just as sometimes one only sees an arc of a rainbow.

I live in the Seattle area, where I have frequently seen small pieces of the 22-degree halo to the sides of the sun, as in Figure 2. Referred to as "sun dogs" (formally called parhelia), these small fragments of bright color seem to just hang in the sky, unrelated to the sun or any other object. Also visible in Figure 2 is the upper tangent arc, so named because it is tangent to the top of the 22-

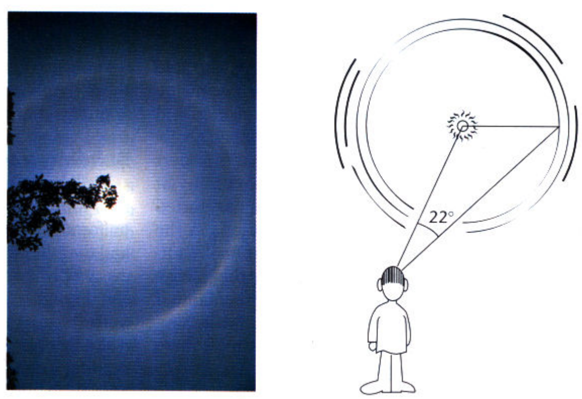

1 The solar halo is a bright circle surrounding the sun.
Photo by Robert Greenler

2 Sun dogs, or parhelia, are small pieces of the solar halo seen off to the sides of the sun. Here the sun dog appears on the horizon, to the right of the sun. The 22-degree halo and upper tangent arc are also visible.
Photo by A. James Mallmann

3 The circumscribed halo occurs when the upper tangent arc (also visible in Figure 2) wraps around the sun and ultimately joins the lower tangent arc. You can see it most clearly at the left and right sides of the 22-degree halo.
Photo by Robert Greenler

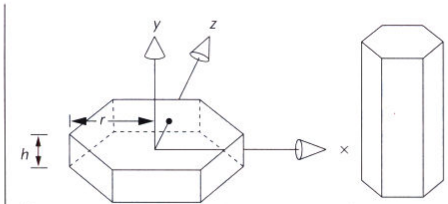

4 Hexagonal ice crystals can take the form of (a) a plate or (b) a pencil.

degree halo. As the sun rises, this upper arc wraps around and ultimately joins the lower tangent arc to form the circumscribed halo, visible in Figure 3. Finally, on rare occasions you can see a dim, large, 46-degree halo centered at the sun. Although you can see some color effects in these phenomena, they are not rainbows. Rainbows are a complex subject, which I leave for another time.

A single, simple mechanism is responsible for all these solar phenomena. Although people have been drawing and hypothesizing about solar halos and sun dogs and such for many years, the pioneer who developed a solid quantitative demonstration of their origins is Robert Greenler, a professor of physics at the University of Wisconsin, Milwaukee. He and his colleagues started publishing the results of their investigations in the 1970s, and he recently wrote a marvelous book summarizing their work.[1] In this issue, I present the essence of Greenler's elegant approach. In my next column, I will show how to generalize it to include full-color simulations, which involves bringing in some tools more familiar to computer graphicists.

Sunlight and ice

Greenler's basic hypothesis is that solar displays arise from the interaction of sunlight with clouds of simple ice particles. We're all familiar with the tendency of ice to crystallize in six-sided forms, as in snowflakes. But the humbler hexagonal prism, shown in Figure 4a, is all we need to bring about the effects discussed above. Actually, we will find it useful to distinguish two forms of the crystal: the plate (in Figure 4a), and the pencil (in Figure 4b). I consider a crystal to be in the pencil form when the ratio of length to radius exceeds about 2. (Toward the end of the column you'll see why I put the dividing line there.) According to Greenler's hypothesis, all the solar phenomena mentioned here result strictly from light reflecting from and refracting through these little hexagonal ice crystals as they take on different shapes and orientations.

To duplicate Greenler's results, we need to make a bunch of simplifying assumptions, many of which will be removed in the next column. Physicists are famous for simplifications. There's an old joke about a physicist who was asked to determine exactly how horses run. His first step was to simplify the problem—he promptly began studying the physics of a spherical horse. But used with restraint, simplifications let us begin our attack on a problem, and if they're well chosen, let us focus on the most important aspects of the problem first, leaving details for later. We will follow that course here.

First, we will ignore reflection altogether, and consider only light that travels through the body of the crystal. Second, we will ignore all effects that could serve to diminish the energy of the light as it interacts with the crystal. That means we'll ignore Lambert's law where the light strikes the crystal, Fresnel's law both where the light enters and leaves the crystal, absorption within the crystal, and polarization. Third, we will ignore all but first-order effects; that is, we will pretend that a given ray of light leaves the sun, interacts with exactly one crystal, and then reaches our eye—none of the light involved in this step ever encounters another crystal. Finally, we will assume a thermally uniform and empty atmosphere, where light travels in straight lines and is never scattered or absorbed by airborne particles. With these simplifications, we can begin to make our first approximation to the solar halo.

Building a solar halo

We'll begin by studying plate crystals with a thickness-to-radius ratio of about 1.5. Assume that they're uniformly distributed in a cloud between us and the sun, and that the wind is kicking up, so they're all tumbled every which way. What would we see if we looked at the sun through a cloud of these crystals?

The most straightforward approach might be to build a 3D model of one crystal, instance it a few tens of thousands of times throughout the cloud, then ray-trace it. This probably would work, but very slowly. When Greenler first studied this problem in the 1970s, the machines were much slower than today's, so he was forced to find a more efficient approach. He invented a technique that is so elegant, I will make it our starting point.

Greenler's idea was to model a single prototype crystal, then give it a random orientation. He finds where the crystal has to be in space to refract light back into our eyes and marks that direction as the source of light energy. Then he returns to the prototype, picks a new random orientation, and finds its new position. Let's look at this in a bit more detail.

First we build a prism with hexagonal top and bottom faces and rectangular sides. I chose to build mine in a left-handed coordinate system, as in Figure 4a. To simulate the wind's effect, we rotate this crystal around all three axes into a random orientation. Then we create a ray of light that starts at the sun and ultimately strikes this crystal. We follow the ray forward as it refracts upon entering the ice and refracts again when leaving. If we're to see this outgoing light, then we must be looking at the crystal along the direction of the outgoing light, as shown in Figure 5. In other words, the final direction is a source of light energy. So we add into an accumulating image a little bit of light coming from that direction. This approach shares a common spirit with Whitted's 1980 approach to ray tracing, where he suggested following rays backward from the eye into the scene.

Figure 6 shows the result of this simulation for 40,000 light rays, crystals with a thickness-to-radius ratio of 1.5, and a sun at the horizon (so technically we would not be able to see the lower half of the figure). My simulations use a viewing angle of 180 degrees, so the entire visible hemisphere is projected into the image (like photographing the world reflected by a shiny sphere, as in Figure 2). In this figure, we can see energy from both the 22-degree halo and the 46-degree halo. Both halos have the right form: an abrupt inner edge at the correct angle, trailing off to the outside. The dots are color-coded following Greenler's suggestion that different transport paths through the crystal give rise to different effects.

Paths to enlightenment

There are five different types of transmission paths through the crystal, illustrated in Figure 7. I call the top and bottom hexagons the lids and the six rectangular faces the sides. If light enters through a lid, it can come out through the other lid (so I call this type of path LL) or through a side (LS). If it enters a side, it can come out through a "direct neighbor," which is a side directly adjacent to the one it entered (SD); a "far neighbor," which is next to a direct neighbor (SF); or the opposite side (SO). These are bi-directional paths; transport from a side to a face is indistinguishable from light going from a face to a side. In the simulations, I color coded the path taken by each ray: LL is red, LS is cyan, SD is green, SF is yellow, and SO is magenta.

Let's think about what these paths represent. Light following an LL path enters a lid, refracts, strikes the other lid, refracts again, and exits in the same direction as it entered. You can prove this to yourself with a little geometry, using the fact that the two lids are parallel. So if we look directly at the sun, we would expect all the LL rays to land right on top of one another, in the center of the sun. In Figure 6, all LL paths are red, and they all land in the center. The same analysis holds for the SO paths, since again we have a ray entering and leaving through parallel faces; the red dot in Figure 6 represents them as well. Paths connecting neighbor sides (SD) can't occur in this kind of crystal. To prove that to yourself, assume an index of refraction of 1.33 for ice and consider that to exit through a direct neighbor, the transmitted ray must make an angle of greater than 60 degrees with the normal of the incident face. (If that doesn't do the trick, think about the critical angle at an air-ice interface.) So LL and SO paths are all the same, and SD paths can't happen, leaving just LS and SF. And as we can see from Figure 6, it does seem that the LS paths (in cyan) are responsible for the 46-degree halo, and the SF paths (in yellow) build the 22-degree halo.

The 46-degree halo in Figure 6 seems less densely populated than the 22-degree halo. In fact, it's far less common to see the larger halo, probably because it's dimmer. How much dimmer, though? Keeping in mind that we decided to ignore energy effects, we can still learn something just by considering how many rays take each of the possible paths. Figure 8 shows the results of sending 20,000 rays through a crystal cloud for differ-

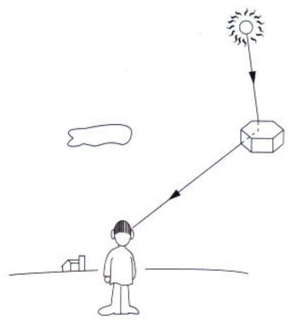

5 To see light leaving a crystal, we must be looking at the crystal along the direction of the outgoing light.

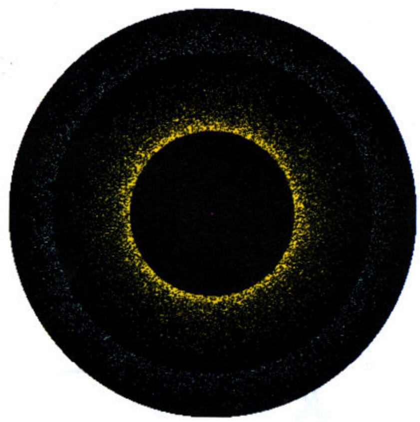

6 The result of the simulation for 40,000 light rays, crystals with a thickness-to-radius ratio of 1.5, and the sun at the horizon.

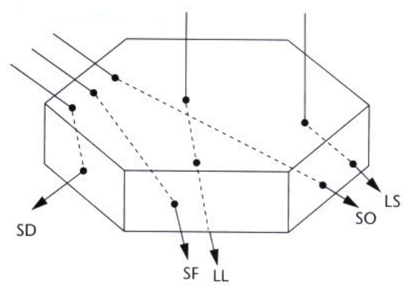

7 The crystal has five different transmission paths for light.

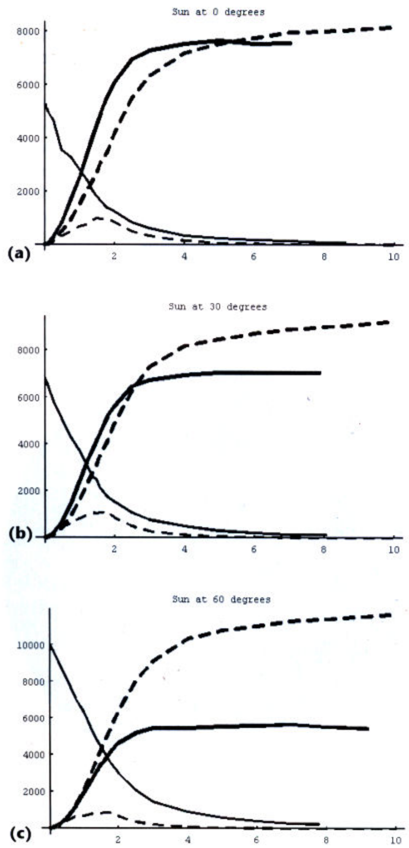

8 **The results of sending 20,000 rays through a crystal cloud for different sizes of randomly oriented plate crystals and different elevations of the sun. (a) The sun at 0 degrees above the horizon, (b) 30 degrees, and (c) 60 degrees. In the graphs the thin solid line=LL, the thin dashed line=LS, the thick solid line=SF, and the thick dashed line=SO.**

ent sizes of randomly oriented crystals and different elevations of the sun. At sunrise (or 0 degrees of elevation), using crystals of thickness 1.5, we find that about 4,600 rays take the SF path and only about 870 take the LS path. Considering also that the LS rays are much more spread apart, it seems that the simulation correctly suggests that the 46-degree halo receives fewer rays per unit visible area than its 22-degree companion.

So far we've assumed random orientations for the crystals because of the presence of a strong wind. But if the wind calms down, then crystals will tend to fall with their hexagonal sides parallel to the ground, like a piece of paper fluttering to the floor. In Figure 9a, I've restricted the orientation of the crystals to only 7 degrees plus or minus on the *X* and *Z* axes (refer to Figure 4), though they are free to rotate about the *Y* axis. The yellow arcs at the left and right are sun dogs. The 46-degree halo here is made up of six different arcs, now distinguished because of the limited crystal orientations. In Figure 9b I've extended the *X* and *Z* rotations to 20 degrees in either direction, and everything begins to spread out and come closer to complete halos.

So far, we've only looked at the sun on the horizon, but it turns out that the phenomena due to plate crystals don't change much as the sun rises.

The plot thickens

Letting the crystals return to random orientations, what happens if we thicken the crystals, turning them from plates to pencils? The graphs of Figure 8 show that for ratios below about 1.5, the LL paths dominate. This makes sense, since the lids are the largest faces in the crystal. Above that ratio, the SF and SO paths become more common, as the side faces grow larger. The LS paths have a small peak around 1.5, so that size best shows the 46-degree halo—that's why I chose a ratio of 1.5 for the earlier figures. If we enlarge the crystals to a ratio of 6, as in Figure 9c, then the number of rays that contribute to the 46-degree halo drops dramatically.

As before, let's now move from a windy to a calm day. Unlike the plate crystals, pencils fall with their rectangular sides facing the ground, like a log in a river. The simulation in Figure 10a shows crystals with a length-to-radius ratio of 8 (the small circle marks the location of the 22-degree halo). They fall with their long axis parallel to the ground, but are free to rotate about the other two axes. Here we see the upper tangent arc effect as in Figure 2. As the sun rises, the upper and lower tangent arcs begin to wrap around the 22-degree halo. Figures 10b, 10c, and 10d show the results with the sun at 11 degrees, 30 degrees, and 50 degrees; as the sun rises, the circumscribed halo comes ever closer to the 22-degree halo.

Implementation and notes

If you're interested in replicating these figures and experimenting with crystal formations, it isn't difficult; the whole program isn't much more than a simple ray tracer. I'll sketch out my implementation and mention a couple of notes.

I place the camera at the origin, looking down the positive *Z* axis, and place the sun at infinity on the positive *Z* axis. The crystal is originally centered at the origin, is assigned an orientation, and then gets moved into position on the unit hemisphere in front of the camera. In these simulations I use a viewing angle of 90 degrees—the most fishy of fish-eye lenses without having eyes on the back of your head! It's a bit extreme, but it allows me to see big effects, like the 46-degree halo (and the light that fades out from it). It also makes it easy to plunk down a dot in the right place; if the normalized direction vector of the ray arriving at the eye is (x, y, z), you can just put a point at $(-x, -y)$. As the sun rises, I like to keep the sun in the center of the image; I just rotate everything down to the horizon before plotting.

The main loop in the program picks an orientation for the crystal, rotates it into place, traces a ray, and plots the ray's outgoing direction. Then I pick a new orientation and repeat. To orient the crystal, I pick some random angles and build a rotation matrix. Because it's pure rotation, I use the same matrix to transform both the points of the crystal and the normals.

To start a ray, I pick a "goal" point. This is a point in a disk of radius 1, centered at the origin, lying in the *X-Y* plane. Then I find the direction back to the sun. Because

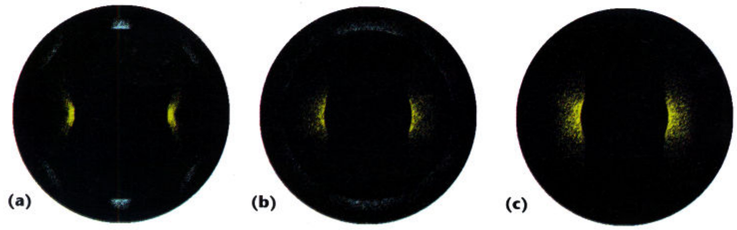

9 Orientation of the crystals is restricted to ±7 degrees on the *X* and *Z* axes (a). The yellow arcs at left and right are sun dogs. The *X* and *Z* rotations have been extended to 20 degrees in either direction (b), resulting in more complete halos. Enlarging the crystals to a ratio of 6 (c) causes the number of rays contributing to the 46-degree halo to drop dramatically.

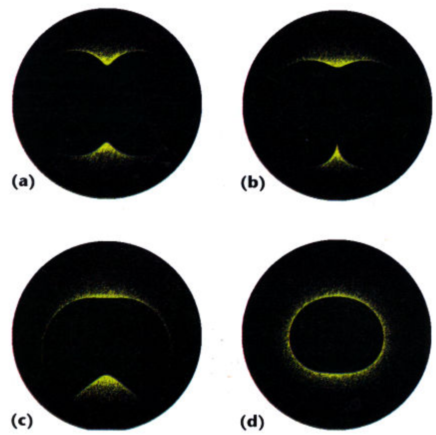

10 The simulation shows crystals with a length-to-radius ratio of 8 and the sun at 22 degrees (a), which fall with their long axis parallel to the ground in the absence of wind. With the sun at 11 degrees (b), we see the arcs begin to wrap around. With the sun at 30 degrees (c) and 50 degrees (d), the circumscribed halo comes closer to the 22-degree halo.

the sun is so far away, I treat it as a point light source generating parallel illumination rays. Since I assume the sun rises in the *Y-Z* plane, the vector pointing from the goal to the sun is $(0, \sin(\theta), \cos(\theta))$, where θ is the angle of elevation of the sun. I push the goal point back along this direction by a length of 2, or twice the thickness, whichever is larger—this is just to make sure that I start outside the crystal. This transformed goal point is the origin of the ray, and the ray's direction is the opposite of the direction in which I pushed it.

The first ray-object test is to find out if the ray strikes the crystal, since the crystal will not occupy the entire unit disk. If the ray misses, I just discard that orientation and start the loop again. Otherwise, I follow the ray through the crystal using standard ray-tracing techniques. It's important to keep the selection of the goal point and the rotation of the crystal decoupled; if they correlate, it will introduce visible artifacts into the results.

The hexagonal prism is a convex solid of convex faces. For convex faces, I've always liked the enclosure method for the point-in-polygon tests. For each edge of each polygon, I build a wall—a plane perpendicular to the plane of the polygon, passing through the edge, and pointing inwards. When I find the intersection of the ray with the plane of the polygon, I test that point against each of the plane equations that make up the walls. If it's on the positive side of all the walls, the point is within the polygon. As soon as it's on the negative side of any wall, you can quit the test.

The only thing left is to pick a wavelength for the simulation, compute the index of refraction of ice at that wavelength, and start orienting crystals.

I built a small pushbutton-style interface that allows me to set global parameters such as the number of rays to trace, the thickness of the crystal, and the angle of the sun. I can select which paths I want to have plotted, the initial orientation of the crystal, and the range of random rotations that can be applied once it's been positioned. These two steps let me place the crystal properly (for example, the plate with hexagon side down, or the pencil with a rectangular side down), then constrain how much the wind can cause it to rotate. When a run is complete, I print out some statistics, such as the number of rays that followed each type of path and the number of points plotted. I also keep track of how many rays underwent total internal reflection (TIR) within the crystal and simply extinguish all such rays.

Next time

In the next column I'll talk about extending this simulation to include energy losses and reflection, creating a smooth image on the screen rather than a cloud of dots, and making it all work in color. In the meantime, you might want to read Greenler's book or do some experimenting on your own. Most of all, observe the skies around you a little more closely. Two other great books on the interplay of light and the atmosphere are Minnaert's classic[2] and the more recent book by Meinel and Meinel.[3] I find that with increased knowledge comes increased pleasure in the beauty of our world. I hope you find the same. ■

References

1. R. Greenler, *Rainbows, Halos, and Glories*, Cambridge University Press, Portchester, New York, 1980.
2. M.G.J. Minnaert, *Light and Color in the Outdoors*, Springer-Verlag, Berlin, 1937 (revised 1985 edition).
3. A. Meinel and M. Meinel, *Sunsets, Twilights, and Evening Skies*, Cambridge University Press, Portchester, New York, 1983.