

FIGURE 5.1 Tangent arc as calculated by Wegener [85] in the 1920's. The responsible crystals are oriented columns, and the ray paths are those that enter and exit alternate prism faces. The sun elevation is 20° . Wegener used stereographic projection from the zenith, which gives the halo a slightly unfamiliar look. The significance of the two red segments is explained on page 52.




FIGURE 5.2 Computer simulations showing halos from oriented columns. The sun elevation and the projection are the same as in Figure 5.1. (Left) Allowing only ray paths for the tangent arc, that is, ray paths that enter and exit alternate prism faces. The tilts of the crystals are zero. (Middle) Allowing all ray paths. The tilts of the crystals are still zero. (Right) Allowing all ray paths. The tilts of the crystals are several degrees.

Halo Simulations

We understand halos far better today than thirty years ago. Several factors are responsible, but none more significant than the simulation of halos by computer. Prior to the computer, although it was quite clear in principle how to predict the halos due to specified crystal shapes and orientations, in practice the calculations were often too cumbersome to be feasible. With the computer the calculations became effortless. The results, moreover, could be displayed graphically in simulations that *looked* like real halo displays.

A comparison of Figures 5.1 and 5.2 gives some suggestion of the power that halo simulations can bring to bear. Figure 5.1 is an example of a theoretical halo prediction made prior to the computer. It is a relatively sophisticated example, taken from Alfred Wegener's elegant *Theorie der Haupthalos* [85] from about 1925. The halo in the figure is the tangent arc, calculated for a sun elevation of 20° . Wegener was able to calculate the entire region of sky occupied by the halo, not just the so-called caustic curve, where the halo tends to be brightest. This depiction was about as much as anybody could hope for at the time.

Figure 5.2 shows three computer simulations analogous to Wegener's diagram. The simulation in the left-hand diagram is the closest to Wegener's diagram; in both diagrams the responsible crystals were oriented columns, the tilts of the crystals were zero, and the allowed ray paths were only those that make the tangent arc. The two diagrams agree with regard to the region of sky occupied by the halos, but the computer simulation enjoys a crucial advantage in that, unlike the Wegener diagram, it indicates intensity variations within the halo. In general, the simulation gives a more realistic and appealing depiction of the halo.

The middle simulation in Figure 5.2 allows all ray paths through oriented columns, not just the ray paths that make the tangent arc. It therefore shows all

FIGURE 5.3 Point (heavy dot) on the celestial sphere in the direction of the crystal axis. The point gives a partial depiction of the crystal orientation. The crystal is at the center of the sphere.





FIGURE 5.4 Depiction of the crystal orientations for the simulation in Figure 1.4. Each of the points shown on the sphere is determined by a crystal orientation as explained in Figure 5.3. (*Left*) For the plate orientations. (*Middle*) For the column orientations. (*Right*) For the random orientations.




FIGURE 5.5 Depiction of the crystal orientations for the simulation in Figure 1.7.

halos (in the given region of sky) due to the crystals under consideration. This was not possible prior to the computer, since there was no way to anticipate all of the ray paths that would be significant. And even for the paths that could be anticipated, it was difficult to judge their relative importance in making the halos. This diagram is a big step forward.

The right-hand simulation is the same except that, instead of being perfectly horizontal, the crystals were given tilts of several degrees. The simulation shows immediately the effect of the larger tilts. While some of this information could also have been obtained from hand calculations, the computer simulation is far more effective in displaying the results.

Seeing crystal orientations

The computer is also good at displaying the crystal orientations that go into a simulation. The resulting picture is usually more enlightening than an analytic description of the orientations. Figure 5.3 shows how a crystal orientation can be partially represented as a point on the celestial sphere. In this way any set of crystal orientations appears as a set of points on the sphere.

Figure 5.4 is the picture for the crystal orientations that went into making the simulation in Figure 1.4. The column orientations appear as a band of dots around the equator, and the plate orientations appear as a concentration of dots centered at the north pole. The random orientations are of course spread uniformly over the sphere.

Figure 5.5 is the same, but for the much better halo display in Figure 1.7. The crystals obviously have much smaller tilts; the dots for the column orientations are barely distinguishable from the equator, and the dots for the plate orientations are clustered tightly about the north pole.

Revealing the ray paths

Figure 5.6 is an all-sky simulation of halos due to oriented columns. The simulation contains 20,000 halo points, each of which arises from a ray path through some crystal. The computer can remember the ray path for each halo point and can be instructed to sort and count ray paths, thus showing which types of ray paths were responsible for the halos. The table in the figure shows the result of sorting the 20,000 ray paths. For example, the tenth line of the table shows that 2.8% of the 20,000 ray paths fell within the broad category of 1352 (Again see Figure 6.1 for face numbering.)

But which ray paths made which halos? We next show by example how the computer can help to answer this question. As we do so, please bear in mind that,

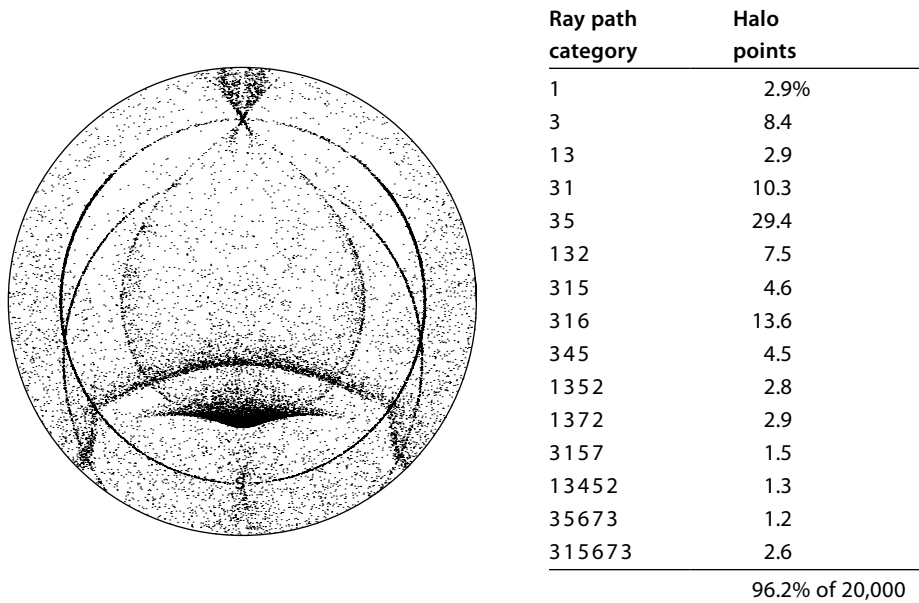


FIGURE 5.6 (Left) All-sky simulation of halos made by oriented columns. There are 20,000 halo points in the simulation. (Right) A sorting of the 20,000 ray paths. Path categories with less than one percent of the total are omitted from the table. The sorting used here is rather coarse, in the sense that the criterion for grouping of ray paths is relatively easy to satisfy; this keeps the table small but hides some information.




FIGURE 5.7 The halo points in Figure 5.6 whose ray paths are in the categories 1352 (left) and 1372 (right).

although we give some details in order to make the example more concrete, the details themselves are unimportant at the moment.

Continuing, then, with the ray paths in the category 1352 mentioned above, we can ask the computer to plot just the halo points that these paths made. The halo points turn out to lie on the left side of the halo known as the subhelic arc. Similarly, the halo points with paths in category 1372 lie on the right side of the subhelic arc. See Figure 5.7.

Thus the ray paths in categories 1352 and 1372 made the subhelic arc. But were there other paths that also contributed? If we were to forget the computer for a moment and just apply the laws of reflection and refraction (page 38) to the paths 1352 and 1372, we would discover a characteristic property of the halo made by these paths: the halo point can always be found by rotating the sun point through an angle of $\pm 120^\circ$ about the crystal axis, with the crystal axis of course being thought of as passing through the center of the celestial sphere. Neither the property itself nor the details of its justification are at all important at the moment; what is important is that there is some criterion—independent of ray path—for deciding which points of the celestial sphere are on the subhelic arc. The computer can therefore be told to sort just those ray paths whose halo points are on the subhelic arc.

The results of such a sorting are shown in the table in Figure 5.8. There were 1,486 halo points on the subhelic arc. The bulk of their ray paths, about 76%,




FIGURE 5.8 (Left) Same as the simulation in Figure 5.6 but only plotting halo points on the subhelic arc locus; the points are chosen according to location, not according to ray path. There are 1486 such points. (Right) Sorting of the 1486 ray paths. Path categories with less than one percent of the total are omitted.

do indeed fall within the two categories 1352 and 1372. There are also some paths—in the first five lines of the table—that appear in the table by accident, in the sense that their halos happen to intersect the subhelic arc; these paths are irrelevant to the subhelic arc. But there are also two lines of the table—the last two—that reveal new ray path categories for the subhelic arc, namely, 134562 and 165432. That these two ray paths are genuine subhelic arc ray paths can be seen either by plotting their halo points—from among the original 20,000 halo points—or by realizing that the two paths always satisfy the geometrical criterion for their halo points to be on the subhelic arc.

The sorting that was used in making the tables in Figures 5.6 and 5.8 was intentionally made to be rather coarse. For example, the sorting considers the two ray paths 1352 and 1382 to be the same and therefore puts them both in the same category, namely, 1352. There is some justification for doing so, since the two internal reflections 3 and 5 have the same effect on the light point as 3 and 8. However, in the first case the two internal reflections are from alternate prism faces, while in the second they are from adjacent prism faces, so that for some purposes you might want to distinguish between the two paths. Moreover, if the computer were asked to take the 20,000 halo points from before and to plot first those that have ray paths like 1352, and then those that have ray paths like 1382, you would find that the two ray paths tend to light different parts of the subhelic arc, so that again there is some reason to distinguish the two paths. When we want to understand a halo in detail, we therefore use a finer sorting than that in the two tables above.

The subhelic arc is an interesting halo. Wegener [85] was aware that the existence of the subhelic arc was a theoretical possibility, and he was able to calculate the expected shape of it using the geometrical characterization mentioned above, but he thought the arc would be too weak to be visible. He was very nearly right; the subhelic arc is a very rare halo. Its existence, however, is no longer in doubt. It is common in the big halo displays in the Antarctic interior (Figures 6.7 and 6.8), and it is seen on rare occasions elsewhere as well (Figure 5.9).

How the simulations are made

The first computer simulations of halos were made by Robert Greenler [21] and his colleagues in the 1970's. Modern simulations are closer to those of Pattloch and Tränkle [56] from the 1980's.

In making a halo simulation the crystal shape is specified in advance. The crystal orientations are specified as well, perhaps with some orientations being more likely than others. The problem, then, is to display the halos that are theoretically expected from the specified crystal shape and orientations.



FIGURE 5.9 Halo display with remarkably strong subhelic arc. Fairbanks, Alaska, March 6, 2002. The responsible crystals were at normal cloud levels, not near the ground.

The solution is much in the spirit of Chapter 4. A sun ray is allowed to fall on a suitably oriented crystal, the ray is followed through the crystal, and a dot is plotted at the halo point. The process is then repeated. After many repetitions, halos begin to appear as concentrations of dots on the celestial sphere.

Of course, all of the preceding is going on mathematically. The crystal is described mathematically, the light rays are described mathematically, and so forth. There are no real crystals, no real light rays.

There are some differences between what the computer does in making the simulations and what was described in Chapter 4. In making the simulations, no ray path is specified in advance. When a ray from the sun encounters a

face of the crystal, it gives rise to a reflected ray outside the crystal and a transmitted ray within the crystal. The computer decides which of the two rays to follow by calculating their intensities and treating the (normalized) intensities as probabilities. If, for example, the reflected and transmitted rays have intensities of .1 and .9, respectively, then ten percent of the time the computer will choose the reflected ray to follow, and ninety percent of the time it will choose the transmitted ray. If the reflected ray is chosen, then its light point is plotted in the simulation. If the transmitted ray is chosen, then it is followed within the crystal until it encounters a second crystal face. Again there will be a transmitted and a reflected ray, and again the computer decides which to follow according to their probabilities. If the transmitted ray (now external) is chosen, then its light point is plotted in the simulation. If the reflected ray (now internal) is chosen, then it is followed to the next crystal face, and so on. As soon as a halo point is plotted, the computer stops and goes on to consider another sun ray and another orientation for the crystal, and the process repeats. When large numbers of dots are plotted, the density of dots in a given region gives an indication of intensity. The simulations show intensity variations within a halo as well as intensity variations from one halo to another.

The computer program that makes the simulations is rather powerful. For example, it can be adapted to simulate halos made by crystals of any polyhedral shape, convex or not, even if the crystals are strongly birefringent. It can therefore simulate halos made by crystals of virtually any mineral, not just ice. The program is not magic, of course; the programmer must still tell the computer what mineral, what crystal shapes, and what crystal orientations to work with.