Re: Ray inside a cone.

Source: http://sci.tech−archive.net/Archive/sci.math/2006−11/msg03977.html

- *From*: "ghoul" <ghoul@xxxxxxxxxxxxx>
- *Date*: 13 Nov 2006 22:51:50 −0800

pierre.bornsztein@xxxxxxxxx wrote:

> "A ray of light is issued from a point interior to a given cone. All the internal surface of this cone can reflect the light.
> Prove that after a finite number of reflections, the ray will never meet the cone again."

There's an easier "kaleidoscope" solution for the 2d case. Say the 2d cone−oid (?) subtends some angle theta. Say K*theta is 180 degrees. Geometrically this corresponds with a 360 degree "fan" being filled with 2K congruent copies of the cone−oid, meeting at a point. Reinterpret the concept of reflection off a boundary line as passing through the boundary line but reflecting the entire diagram. This is justified by the reflective symmetry of the diagram in any boundary, as well as the usual description of optical reflections (angle of incidence = angle of reflection).

Now note that a line drawn through this diagram crosses at most K−1 lines (or is coincident with two of them), where each crossing corresponds with a reflection in the original problem. Thus, if the initial ray is not coincident with an edge, there are at most ceiling(180 degrees / theta) crossings (i.e., reflections).

Now for the hard part −− three dimensions. We can't take the 2d case forward directly, but we can salvage most of it. Say we start with a cone containing the ray's origin. When we reach the surface of the cone, instead of reflecting, "draw" another cone on the other side, coincident with the first cone along the line connecting the origin and the ray's intersection with the cone(s). Since the tangent plane of the first cone at that point coincides with the tangent plane of the second cone at that point, we can say that reflecting the cone in this way is equivalent to reflecting the ray. Repeated reflections of the ray are thus modeled as repeated reflections of the cone along a single ray.

Now the fun begins. Slice the diagram along the plane defined by the cone's origin and the single ray. Two consecutive "cone−oids" must subtend the same angle because they're mirror images of each other in three dimensions, specifically mirrored in a plane perpendicular to the ray. By induction all cone−oids must subtend the same angle, so we're almost to the 2d case. The only remaining problem is proving that the cone−oids in the 3d case aren't arbitrarily more acute than the cones (which might permit the million−bounce trick−shots even in an obtuse cone).

At this point I don't know how to limit acuteness of the cone−oids relative to the cone. If it's possible, then there is a bound on the number of bounces that depends only on the cone angle, not on the ray's origin. If it's not possible then the "finite number of reflections" would still hold for a ray placed any particular finite distance from the surface of the cone, but there would be no finite bound based only on the geometry of the cone (independent of the ray's origin).

My guess is that there is a way of tying the axial and radial changes together on each bounce to force the light quickly away from the apex.

Maybe we could sandwich any cone between two K−pyramids of "similar" angle, then take K −−> infinity. I'm not convinced the K−pyramid variation is actually easier to prove, however.

Any other ideas? How about for hypercones in 4 dimensions?

.