Chapter 4

Check the lights

Headlights, LED spots and other optical systems produce complicated light patterns. Philips Lighting checks the output of their new designs by tracing millions of rays through the system. A reliable method, but very time-consuming. For advanced solutions, the calculations may take days. Can the study group come up with a faster approach?

Philips Lighting develops optical systems. Since the introduction of LED these systems have become more and more intricate. A naive method to predict the brightness and intensity of a new system is to trace a random ray of light through it. For a single ray this is an easy calculation, but for a realistic image you need to track up to a million rays. The devil is in the details, skipping a small number of rays may result in an incorrect pattern. Altogether the computations can take an entire weekend to run. One of Philips Lightning's long-term-goals is developing faster methods, so they can more easily test and optimize their designs.

Wilbert IJzerman, department head of Philips' LED Platform development group, knew about a different approach, but he never had any time to look into it. When he was home with the flu he started thinking about it again: "The walls were closing in on me and I started tinkering." The key idea is to look at groups of rays that leave the light source close to each other with a similar angle. These rays will roughly follow the same path through the optical system. The initial position and angle of a ray form a two-dimensional phase-space. How can this phase-space be partitioned in areas that show the same behavior? More specifically, how can the edge between two such areas be determined? These were the questions IJzerman and his colleagues asked the study group.

They knew that their problem was tough, so they suggested that the study group focused on two-dimensional models. Mathematician Maxim Hendriks from Eindhoven

University of Technology: “The problem itself was immediately clear, but only later we understood what made it so hard. It is impossible to find exact solutions for some of the problems. You can do it numerically by solving the same problem many times. Actually many, many times. Which is troublesome in practice too.”

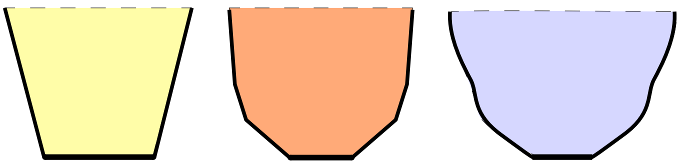

Three cups that are increasingly hard to analyze. The two-faceted cup on the left is the easiest, the multi-faceted in the middle a bit harder and the smooth cup on the right is the hardest.

A cup with two facets

The study group decided to start with a very simple light fixture: a symmetric cup with a flat base and two inclined facets. The entire base is the light source. This models a simple torch or bike light. The light intensity is not equally strong in all directions; it varies with the cosine of the angle of the light ray. These sources are called Lambertian and to the human eye they appear to have the same brightness when viewed from different angles.

The quality of the light bundle is assessed on a screen that is parallel to the light source. In most real life applications it is unknown how far away a wall will be from the lamp. Therefore only the far-field is considered, this may be seen as a target screen at infinity. The distance is in practice always much larger than the size of the optics and hence you can apply the far field approximation.

For an elementary two-faceted cup it was possible to analytically determine the number of reflections before a ray left the cup. To ease counting the reflections, the mathematicians mirrored the cup and not the rays.

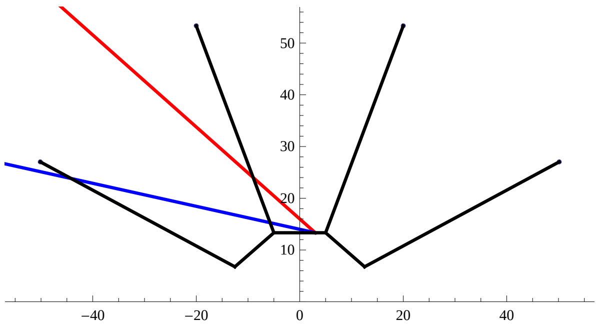

The number of reflections of a ray is the same as the number of reflected cups it passes. So in this example the red ray is reflected once and the blue ray twice.

They divided the phase-space in regions where rays had the same number of reflections. The boundaries between these regions turned out to be nearly straight lines for their simple example. The next step was to calculate how a reflection changes the angle of a light ray and use that to determine the exiting angle of each ray. Finally, from all this they computed the intensity pattern at the far field. For some cups a familiar pattern arose. The light was brightest right in front of the fixture, but just right and left of this bright core was a darker ring, which in turn was surrounded by a bright ring. Hendriks: “I remember this effect from when I held a torch light as a kid. There was always a black spot close to the middle of the light bundle.”

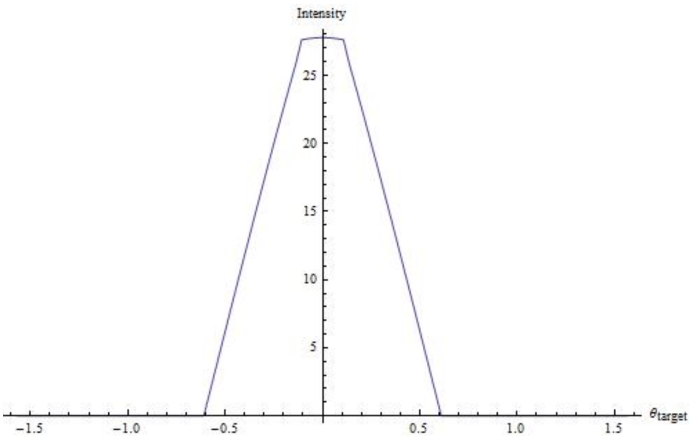

For the example cup suggested by Philips the intensity became a nice and simple peek.

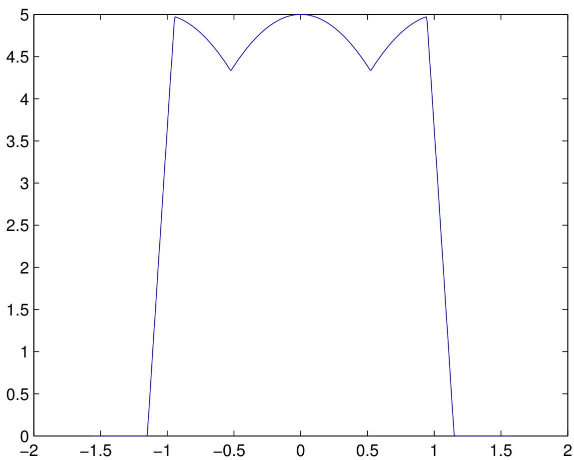

For a cup with slightly different inclined edges the intensity pattern was something like a batman mask. This is the torch light effect.

Different cups

Multi-faceted cups can be seen as a stack of two-faceted cups. By careful retracing the beams of light that emanate from the cup back to the source, the results from two-faceted cups can be used to find the intensity for these multi-faceted cups. The study group also managed to generalize part of the results to polygonal cups.

For smooth cups things became much harder. To make things workable the study group made the rather limiting assumption that the ray trajectory always alternated between the left and right side of the cup. So cups that would reflect a ray from the left side back onto the left side were not allowed. Even with this restriction it was impossible to reconstruct the exact partitioning of the phase-space. The study group showed how finding the boundary of regions in the partition is equivalent to finding the zeroes of a (sophisticated) function. They indicated a whole ensemble of numerical tricks that could do this.

Faster simulations

Finally, the group considered how Philips' current method of ray tracing could be sped up using the partitioned phase-space. Instead of using random rays, one could smartly select a small number of rays to trace and derive the rest of the paths from there. The study group proposed two different methods. The first assumes that the partition of the phase-space is known. The method starts by taking a small preliminary sample of uniformly distributed points from the phase-space. For each of these points their path is determined by traditional ray-tracing. Then a much larger sample of points is randomly selected from the phase-space as an estimation sample. For these points there is no need for ray-tracing, because their path can be computed from the paths of the nearest points in the preliminary sample.

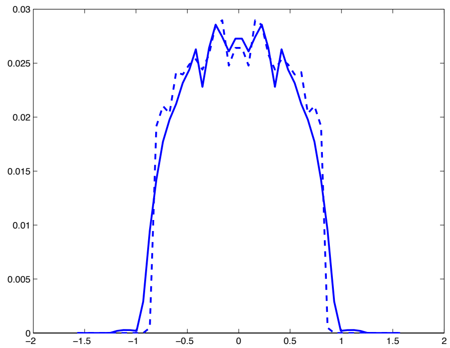

The results from the first method on a smooth cup (solid line). The preliminary sample contained 512 rays, the estimation sample 2^{18}. They compared the pattern with the naive ray-tracing of 2^{18} rays (dashed line).

The resulting intensity pattern was very close to the true profile, but there were some striking differences. This was mainly caused by errors in the approximation of the phase-space partition. Therefore the team came up with a second method that does not need the partition of the phase-space and only makes some very weak assumptions about this partition. As before the method starts by tracing a small preliminary sample of rays and selecting a larger estimation sample. But now for each point in the estimation sample the method takes different groups of nearby points in the preliminary sample that all have the same number of reflections. If the estimation point fits nicely into one of these groups, it must have the same number of reflections as these points. Otherwise, ray-tracing is performed and this ray is added to the preliminary sample as extra data. This method simultaneously estimates the partition of the phase-space and the intensity profile. In general only points close to the boundary of a region in the phase-space need to be traced. The method was tested on a smooth cup with a preliminary sample of 512 points and an estimation sample of 2^{19} points. The method only used 38,006 ray tracings, which is about 7% of the rays. The resulting intensity plot is almost identical to the one from doing 2^{18} ray tracings. This naive method however takes hours, while the new method runs in a few minutes.

First steps

Wilbert IJzerman is not sure that Philips will implement this faster method: “We have to be really sure that we are not missing some details.” He is however very happy with the results from the study group. “For the industry this week is the perfect way to explore new directions. If the study group solves your problem, then it is too easy and not very interesting for further research. If they do not make any progress, your problem is too hard. But if they manage to make some first steps in a week, you get the feeling that you could really do something with this problem if you worked on it for a year.”

Two ideas from the study group were new to him. The first was the backward tracing of the light that was used in the multi-faceted cups: “We had never done something like that in this context. This method gives our designers a better understanding of what happens with the phase-space.” The other fresh idea was searching the boundaries in the phase-space by determining zeroes of a function.

IJzerman has decided that his questions deserve further research, starting with a student and hopefully later a research project from the Dutch Technology Foundation STW. “These are deep questions, both analytically and numerically. I think that the mathematicians in the study group saw that our problems are just as hard as those in academics.”

Team Philips

Jan Bouwe van den Berg (University Amsterdam), Rui Castro (Eindhoven University of Technology), Jan Draisma (Eindhoven University of Technology), Joep Evers (Eindhoven University of Technology), Maxim Hendriks (Eindhoven University of Technology), Oleh Krehel (Universität Erlangen-Nürnberg), Ivan Kryven (University of Amsterdam), Karin Mora (University of Bath), Botond Szabó (Eindhoven University of Technology), Piotr Zwiernik (Eindhoven University of Technology).