Application in Illumination Design

Analyzing LiDAR Return Signal Strengths for Target Optical Surfaces and Atmospheric Conditions

Author

Jake Jacobsen, Ph.D.
Senior Technical Marketing Manager,
Synopsys Optical Solutions

Introduction

Body Light Detection and Ranging (LiDAR) systems are becoming increasingly important with the development and deployment of autonomous vehicles. Pulsed LiDAR is the most common form of automotive LiDAR system, though there are other types being developed such as frequency-modulated continuous-wave LiDAR systems. Pulsed LiDARs generate short laser pulses in the near infrared and then measure the time it takes for a reflected signal to return to a detector. This time-of-flight measurement can then be used to calculate the distance to the reflecting, or target, object. By scanning the laser pulses over the instrument's field of view, a three-dimensional image of the surrounding objects can be generated.

LiDAR systems are complex, dynamic devices involving both optical, mechanical, and electronic components that must work together while interacting reliably with a widely varying external environment. It is the optical part of the system that interacts with the environment, and so it is through optical modeling that we can gain a crucial insight into expected performance under varying conditions. In this paper, we will take a brief look at how illumination ray tracing software can be used by engineers to gain these insights into expected system performance before building expensive prototype systems.

LightTools® illumination design software from Synopsys was used to generate the data and images shown in this paper. The model used to obtain the data shown here is that of a somewhat simplified, spinning, roof-mounted LiDAR system shown in Figure 1. The system is designed to use a launch laser with a wavelength of 905 nm. A tilting MEMs mirror is used to achieve the vertical scanning of the laser over the vertical field of view. An eight-element array of avalanche photodiodes, fed by a four-element lens with an aperture of 21 mm diameter, is used for the detector. The array is oriented to cover the entire 16° vertical field of view. The entire body of the LiDAR then spins 360° to accomplish the horizontal scanning. To simplify our analysis, we have chosen to set the MEMs mirror so that the outgoing light is horizontal. The target surface is a simple flat surface that can be moved and tilted as needed, though more complex targets can certainly be modeled.

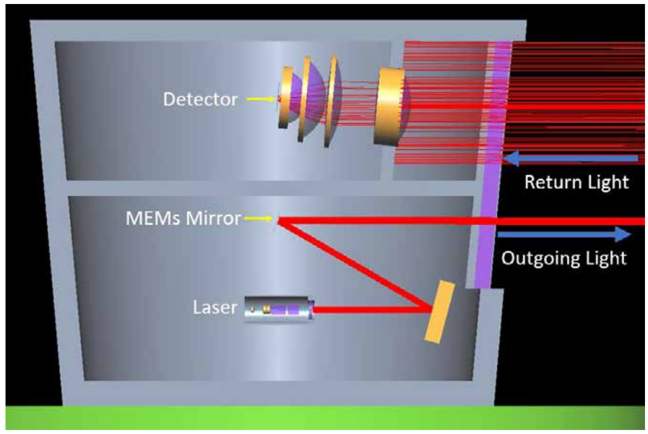

Figure 1: Cutaway diagram of the test LiDAR system

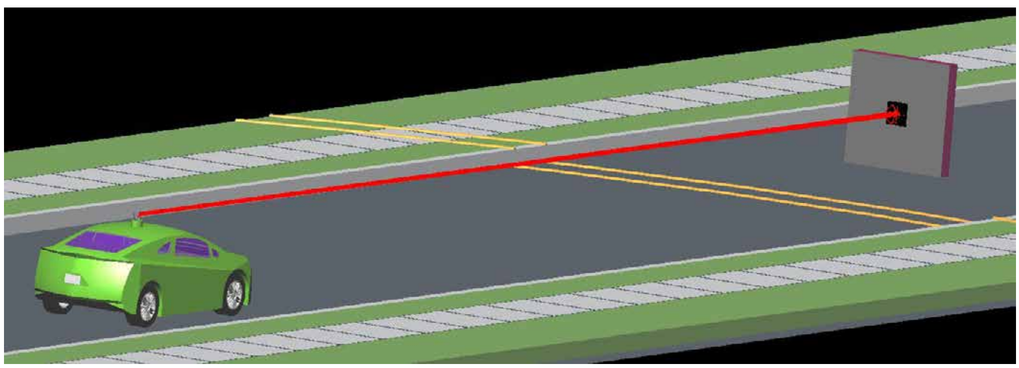

Figure 2: LiDAR system shown illuminating a target cube at 25 m distance

Return Signal Strength

For a LiDAR system to work properly, the return signal must be detectable and distinguishable from any background signals or noise. Sources of background signal can be varied and include sunlight reflecting off the target surfaces or scattered into the receiver from the atmosphere, or even intercepted signals from other, nearby, LiDAR systems. Sources of random noise can also be varied, but often include random noise from the detector and electronics.

The detectors themselves must be very fast and sensitive at the laser wavelength. The typical wavelength chosen for autonomous vehicle LiDAR systems is 905 nm, though some systems make use of 1550 nm laser sources. 1550 nm has the distinct advantage of being more eye-safe, thereby allowing greater output power levels, and having less ambient background from sunlight. However, the reduced availability of fast detectors that are sensitive at the longer wavelength is a difficulty.

Working detection ranges for autonomous vehicle LiDAR systems run from a few meters to between 100 m to 200 m. The launched laser beam impinges on a surface, such as a vehicle or pedestrian, and the light is then scattered back towards the LiDAR system. A very small fraction of that light is scattered into a detector aperture and results in the detected signal. The optical properties of the scattering surface can vary widely and have a direct, and sometimes complex effect on the return signal strength.

Reflective scattering surfaces have both an overall reflectivity, called the Total Integrated Scatter (TIS), and a scatter pattern called the Bi-Directional Reflective Distribution Function (BRDF). Both are dependent on the incidence angle of the impinging light. Together the TIS, BRDF, and incident angle determine what percent of the incident light is reflected into the detector aperture. The BRDF can be quite complex for any given surface, requiring a measurement to accurately determine the surface property. Since the types of surfaces that will be encountered in the real world are greatly varied, we can't obtain direct measurements for all of them. However, we can make some very good progress in our understanding by applying illumination design software to the problem and using some simple scatter patterns at various incidence angles.

If we assume that the target surface is Lambertian, then we can directly calculate the expected return by the following equation:

$$I_{Return} = I_{Launch} * Transmission * P * TIS * \frac{PSA_{Detector}}{\pi} * \cos(\beta)$$

Where IReturn is the return power at the detector; I_{Launch} is the original launch power; Transmission is the transmission of the system including the atmosphere for a round trip; P is the percentage of the beam energy that falls on the target surface; TIS is the total integrated scatter of the target surface; $PSA_{Detector}$ is the projected solid angle of the entrance aperture of the detector lens system as measured from the target surface; and b is the angle of incidence on the target surface. This does not include any detector sensitivity values that you could also factor in.

If we have a circular aperture for the detector system that is oriented normal to the incoming light, we can expand the above equation to:

$$I_{Return} = I_{Launch} * Transmission * P * TIS * \sin^2\left(\tan^{-1}\left(\frac{D_{Detector}}{2 * r}\right)\right) * \cos(\beta)$$

Where $D_{Detector}$ is the entrance diameter of the detector system and r is the distance between the LiDAR system and the target.

As an example, if we assume a launch power of 100 W, transmission and TIS of 100%, 100% of the beam power falling on the target, a detector entrance diameter of 25 mm, and a target at 50 m distance and normal incidence, then we calculate a return power of 6.25 μW. This is a reduction of more than 7 orders of magnitude over the launch power and can be easily confirmed using LightTools (see Figure 3).

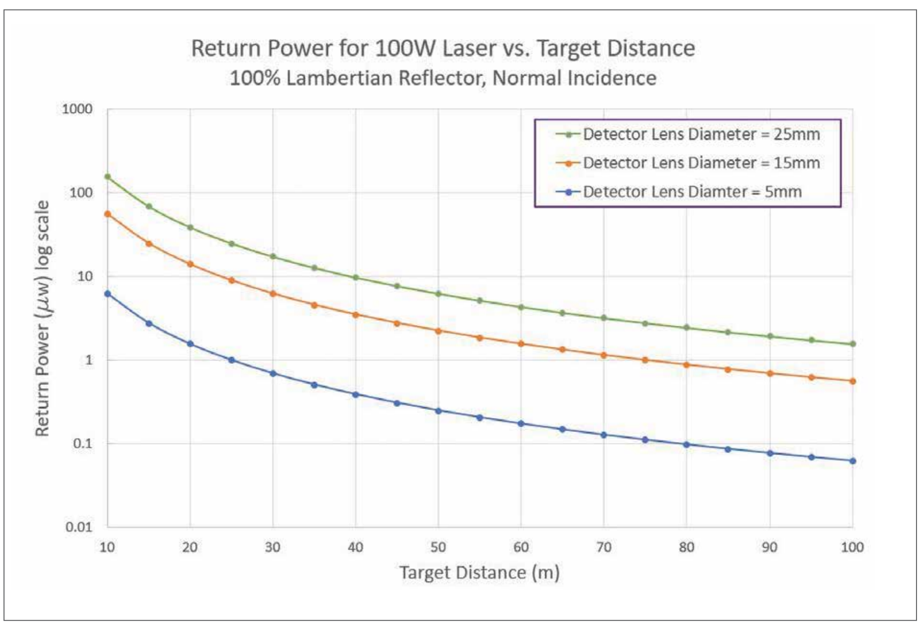

Figure 3: Set of return curves for different detector apertures as a function of target distance using a 100W laser. The target was a 100% Lambertian reflector at normal incidence.

Lambertian surfaces create wide scattered light distributions with intensity patterns that fall off as the cosine of the angle of incidence. Light reflected from Lambertian surfaces looks uniformly bright no matter what the viewing angle is. This makes them reasonable approximations for surfaces that appear matte to the eye such as skin, clothing, and plants when viewed from a distance. Smooth metallic and glass surfaces on vehicles generally do not have Lambertian properties, but exhibit more complex scattering behavior. While for some such surfaces it may be possible to calculate the response, it is much easier to model the surface in LightTools and obtain the result directly, as shown in Figure 4.

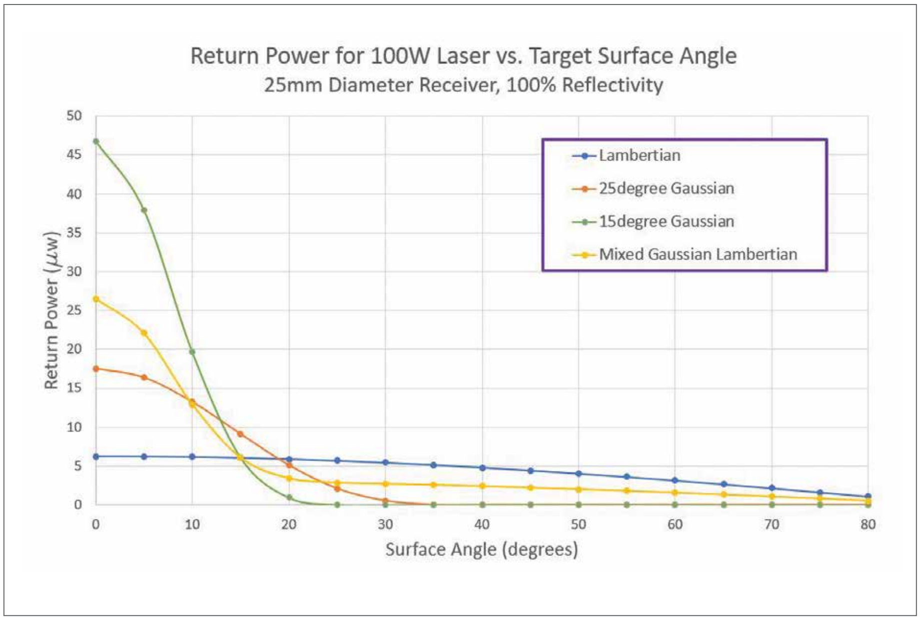

Figure 4: Results from several different scattering surface types for various incidence angles

For all data sets, we use a 25 mm diameter aperture, a 100 W laser and a target distance of 50 m. Figure 4 shows the results from a Lambertian surface, two Gaussian surfaces with differing half-widths, and one mixed surface with a 15° Gaussian (50%) and a Lambertian (50%) component. It is noteworthy that the Gaussian surfaces produce greatly enhanced signals but only at near normal incidence. More complex surface scattering types including measured BRDF can be easily modeled in the same way.

Atmospheric Effects

Another important effect that will influence the return signal is atmospheric conditions such as rain or fog. Raindrops that fall inside the beam will deflect the light that passes through them, excluding some of that power from the return signal. Of course, the system is double-pass, so the precipitation can affect rays going in either direction. The impact of precipitation on the return signal is not easy to calculate directly. The best way to obtain meaningful information is through modeling.

Let's first look at rain and then we will discuss fog as a special case. In order to properly model rain drops in the LiDAR beam, you need to know the size of the drops and the overall density. You also need the index of water at your laser wavelength, which is reasonably easy to obtain. These density and size parameters will vary depending on the strength of the rain. We will use three categories of rain (light, medium, and heavy) with typical drop sizes and densities for each category. Table 1 lists the values that we used for the simulations shown.

Precipitation Condition	Drop radius (mm)	Terminal velocity (m/s)	Drop density (M^-3)	Mean free path (mm)
Light rain (1 mm/hour)	0.5	4.03	131.64	4.84×10^5
Medium rain (6.35 mm/hour)	1	6.49	64.88	2.45×10^6
Heavy rain (24.4 mm/hour)	1.5	7.8	64.98	1.11×10^5
Moderate fog	0.007	-	5×10^7	6.18x104

Table 1: Data used to characterize the three classes of rain and fog used in our simulations

In LightTools, there are at least two methods for simulating the rain drops in the beam. The first method is to create a volume scattering material with the appropriate mean free path between the particles and the correct scatter profile for each. The rain drops do not physically exist in the model, but the rays randomly strike scatter centers in the air between the LiDAR and the target and are scattered appropriately.

The second method, and the one that we utilized for this data, is to use the 3D texture feature with library textures. This feature allows you to create spherical textures to represent individual drops and then translate them up off the base surface into the path of the beam using the z-offset parameter. A simple macro randomly distributes the drops inside a pre-defined volume surrounding the outgoing and incoming beam paths (see Figure 5). You can also use a spreadsheet to calculate the random values and copy them into LightTools. The macro can then be extended to make multiple random distributions, run the simulation for each, and collect the performance data to obtain an ensemble average over a large sample set.

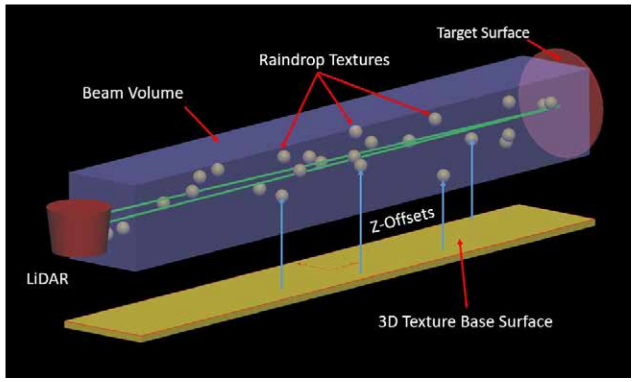

Figure 5: Conceptual layout for the 3D Texture raindrop method. Spherical textures are defined on the base surface and then offset into a volume that encompasses both the incoming and outgoing beams.

The effect of rain drops can be easily seen in the pattern projected onto the target in Figure 6. For each of these simulations we used a 2 mm diameter launch beam with a 0.1° beam divergence (half-angle) and a 50 m target distance. To make the effect of the drops easier to see, we have made the distribution of the source a circular top-hat rather than the more typical Gaussian spatial distribution or even rectangular striped distributions found with multi-cavity lasers. For real system design work, you can model complex source distributions that more closely match the actual laser source distributions.

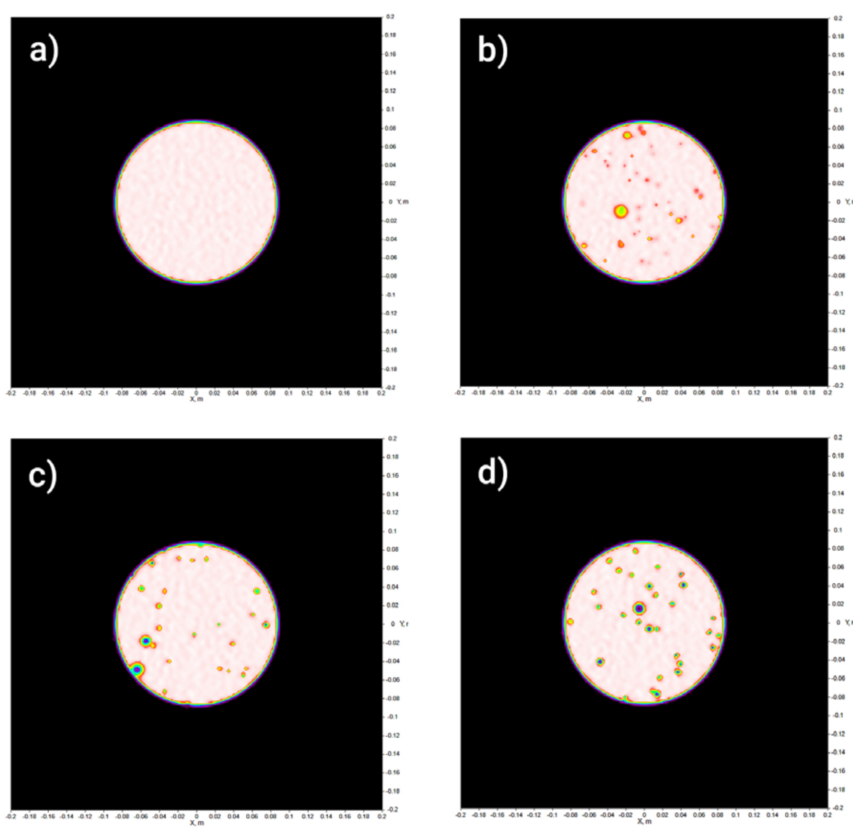

Figure 6: These figures show the illuminance pattern on the target surface in various atmospheric conditions: a) no precipitation; b) light rain; c) medium rain; and d) heavy rain. With heavier rain, fewer drop marks appear, but the average effect of each individual drop is larger.

For each of the different conditions, the size of the drops was held uniform throughout the volume, though this is not a requirement of the software. However, you will notice that the marks from each drop vary considerably. This is the result of their varying position along the length of the beam. Drops that fall in the outgoing beam near the LiDAR system where the beam is small have a greater effect than ones that fall in the beam nearer the target where the beam is much larger. This can be seen in Figure 7, where we happened to have a small drop fall in the outgoing beam near the LiDAR system.

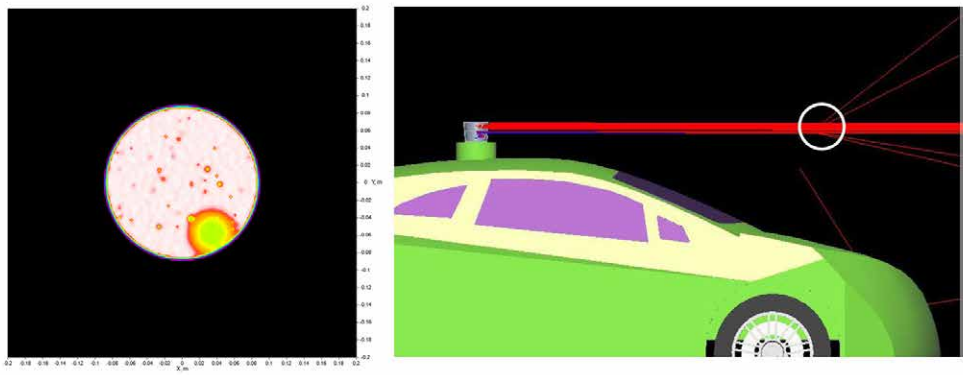

Figure 7: In this figure, we see an illuminance pattern on the target surface with a raindrop falling in the outgoing beam near the LiDAR system, causing an outsized effect on the pattern. The right-hand image indicates the position of the drop (white circle) as evidenced by the scattered red rays.

We have not yet discussed the effect of precipitation on the exit window itself. For spinning LiDARs, such as the one modeled here, this may not be a major issue since the centripetal force may keep any sizeable drops from remaining on the window. However, for fixed systems, the presence of drops on the exit window can be catastrophic, blinding the system to part of its field of view.

Modeling fog is a little different from modeling rain. The primary difference is in the drop size and density. As shown in Table 1, drop sizes for fog are very much smaller than that for rain, and particle densities are much higher leading to many more particles in a given volume. Fortunately, the fog droplet sizes are small enough that we can use Mie theory to model their effect on light. To model fog, we used a volume scattering material with an appropriately sized Mie particle at the correct density for moderate fog (see Table 1). The result is significantly more pronounced than with rain. Figure 8 shows the results of a ray trace using a foggy atmosphere.

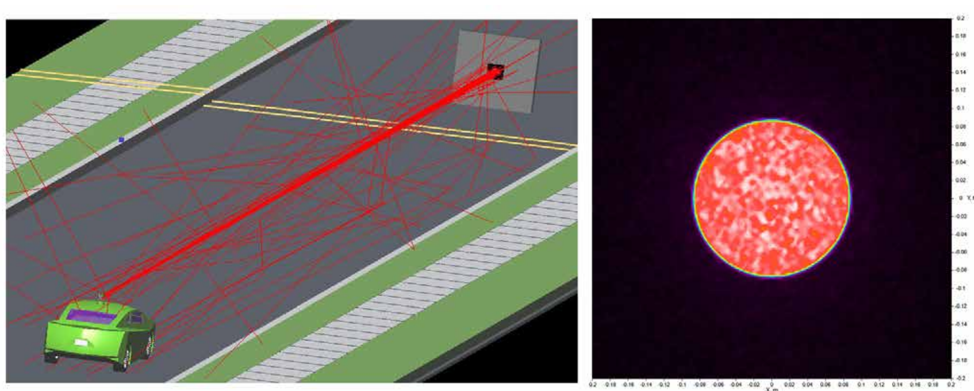

Figure 8: The image on the left shows the LiDAR beam in the presence of a moderate fog. Many rays have been scattered out of the beam by the fog particles. The image on the right shows the illuminance pattern on the target surface. The variation is primarily the result of the random scattering.

While variations in the target illuminance pattern are informative, what really matters to the LiDAR system is the return flux from the target. In this, the results are somewhat surprising. As can be seen in Figure 9, the return signal for the tested rain conditions did not vary significantly from the baseline of no precipitation. Moderate fog, on the other hand, led to a significant signal drop that worsened as the distance increased.

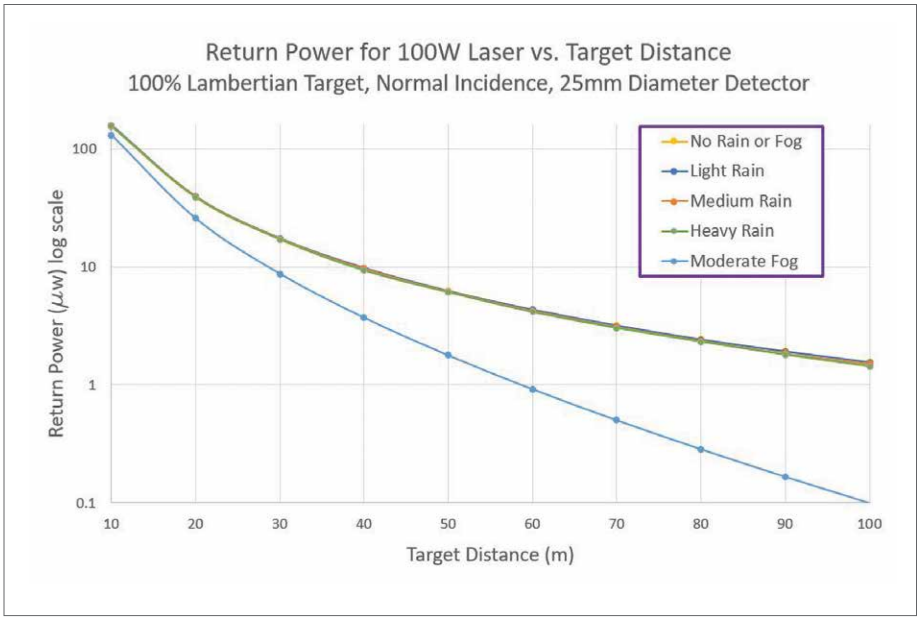

Figure 9: The return power from a Lambertian target is shown at normal incidence as a function of distance in various atmospheric conditions. It is notable that over the distance measured, none of the rain conditions significantly degraded the return signal. However, the fog shows greater than an order of magnitude degradation near the end of the distance range.

Summary

While this paper has not covered the full breadth of LiDAR performance parameters that can be explored with illumination design software, we have sought to show how software can be used to simulate the performance of automotive-based LiDAR systems. Specifically, we have explored the effects on the expected return signal of scattering properties, incidence angles, and atmospheric conditions. Other complex effects can be modeled with illumination software to obtain a better understanding of the expected optical performance of the system under varied conditions. These include, but are by no means limited to, solar radiation interference with the return signal, stray light both inside the LiDAR module itself and as backscatter from precipitation, and the effect of diffraction on the beam shape and size.

To Learn More

For more information about LightTools, visit www.synopsys.com/optical-solutions/lighttools.html. You can also contact Synopsys at optics@synopsys.com to request more information about LightTools and our LiDAR modeling capabilities.

References

http://www.shorstmeyer.com/wxfaqs/float/rdtable.html

05/14/20.CS503583092-OSG WP _ LightTools LiDAR WP.

Pub: May. 2020