

# Seeing Farther with LIDAR Using Tracking

Atulya Yellepeddi, Research Scientist

## Introduction

Technology has always been a tool to free us from our most boring tasks. For many of us in the modern world, few things are more tiresome than being stuck in traffic on our morning commutes, or dealing with highway hypnosis and long weekend traffic for hours on end. While this has fueled much excitement over autonomous vehicles (AVs), the prospect of 2-ton pieces of metal zooming around unattended has also led to a renewed focus on technology to enable their safe operation.

In order to achieve superhuman safety<sup>1</sup>, a detailed 3D map of various dynamic objects (such as other cars, pedestrians, and bicycles) is generally considered essential to an AV. Light detection and ranging (LIDAR) sensors are often considered among the most useful systems to have on board, as they are capable of forming such detailed maps. An example of such a map is shown in Figure 1.

The farther away a self-driving car can reliably detect the presence of an object on the road, the easier an evasive maneuver becomes. Researchers based at Analog Garage (Analog Devices' technology center) have investigated extending the detection range of LIDAR systems and developed a way of using physical constraints on the movement of objects to extend range. To understand this, we first explain the principle of operation of LIDAR.

## How LIDAR Works

LIDAR systems fire a laser pulse at an object and measure the time it takes for the light to reflect off the object and return to the sensor, as shown in Figure 2. By scanning the laser along the horizontal and vertical directions, a LIDAR system forms a complete 3D map of the scene in front of it. Every such map is called a frame. Modern LIDAR systems typically have frame rates between 10 and 30 frames per second (fps).




Figure 1. LIDAR map.

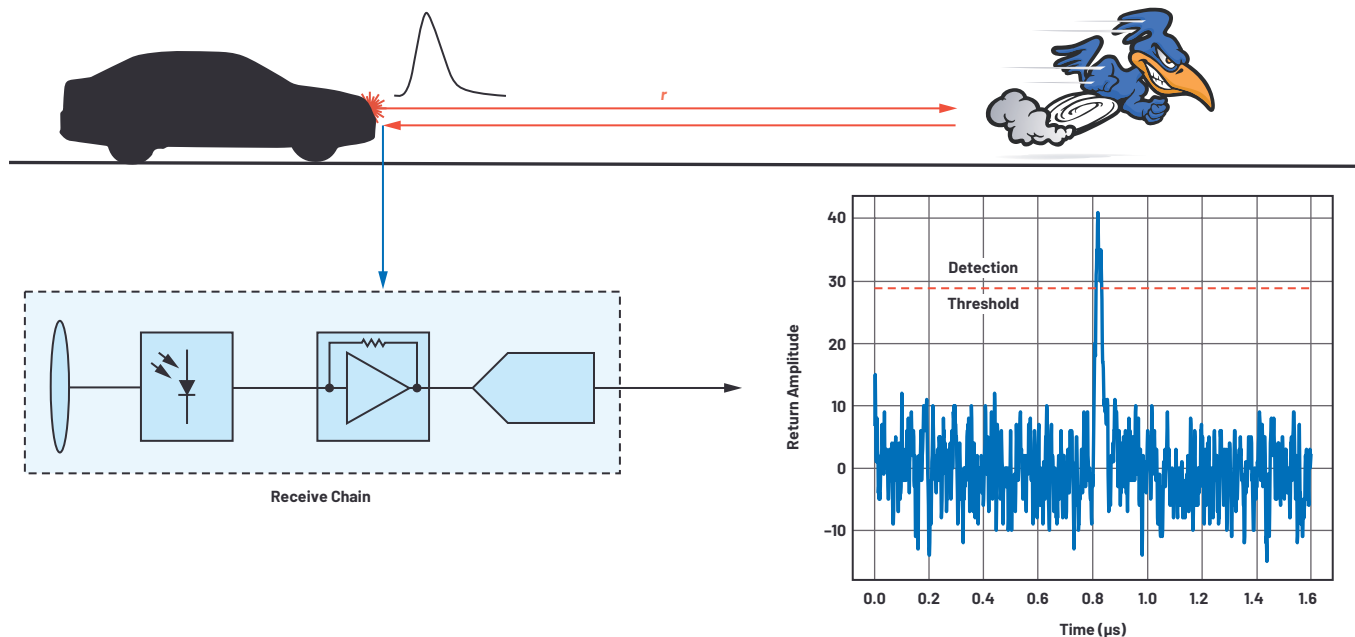


Figure 2. Principle of operation of LIDAR.

<sup>1</sup> While there's no accepted standard for target safety, between 94% and 98% of accidents can be attributed, at least in part, to human error—see the white paper “Safety First for Automated Driving” (SaFAD) at [daimler.com](http://daimler.com).




Figure 3. Effect of distance on LIDAR return signal.

Once the laser has fired in a particular direction, a sensor records the light coming from that direction, converts it into an electrical signal, and searches that signal for the location of the laser pulse shape, using a technique called a matched filter. The output of the matched filter is compared to a threshold, and if the signal crosses the threshold, a detection is declared.

## Detection Limits

Of course, nothing in the real world is ideal. The detection process in LIDAR introduces noise—both electrical noise from the various components of the receiver and optical noise from the detector itself. Thus, an object is detectable only so long as enough light is received from it that the matched filter is able to distinguish that signal from noise.

Physics tells us that the intensity of light falls off as the square of the distance that it has traveled from its source. Practically, that means that the amount of reflected laser light received by the LIDAR system from an object 200 m away is only a quarter of the light that would have been received had the same object been 100 m away.

It follows that our matched filter will have a harder time seeing objects farther away, and, in the extreme, when an object gets far enough away, it will become invisible to the LIDAR system. Figure 3 demonstrates this—the return signal amplitude from the car we are imaging falls off precipitously with the distance, and at 220 m away, the signal is essentially indistinguishable from noise and is missed by the detection threshold set.

We could get around this problem by setting a really low threshold of detection, so that the car at 220 m in Figure 3 is visible. Obviously, given the SNR level, we'd also detect a lot of noise. Now, a full 3D frame of data has a number of flashes in it—some of which correspond to an object, and some of which are just noise. As an example, Figure 4 shows all the detections (post-thresholding) in just one vertical slice (that is, at a fixed vertical angle) of a LIDAR frame. Most of the detections are just noise, but some do correspond to a real object. How can we know which is which? While this is difficult to do with just one frame, it becomes more feasible after we've seen a few frames of data.

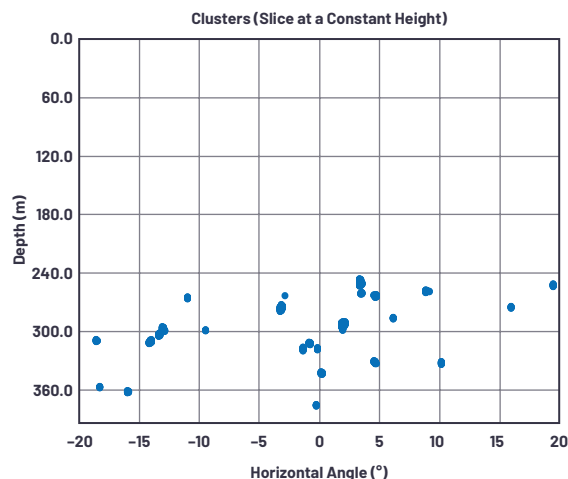


Figure 4. Single vertical slice of LIDAR frame.

## The Firefly Process

To understand why, we can model the flashes like this: suppose there's a firefly buzzing around a box, and we see a flash from the firefly at regular intervals. Unfortunately, we also see random flashes from the environment, and those could occur anywhere. To make matters worse, we sometimes miss flashes from the firefly, and the position we measure for the firefly is usually not quite perfect.

The fundamental question we ask is "Given a sequence of flashes, where each one is a flash from a single frame, can we tell whether the entire sequence came from a firefly or not?" The technical term given to such questions is hypothesis testing. The information we have to make a decision is that the frames arrive 10 times per second (for a frame rate of 10 fps), and a firefly can only move in a physically reasonable way in that time. For instance, a firefly cannot travel the length of the box in a frame, as that would be a physically unrealistic velocity; and it cannot reverse direction in 2 frames, as that would be a physically unrealistic acceleration.

Put a different way, the information that we can use is that the track followed by a firefly must be a track that a physical object could, indeed, have taken. Applying these track-based constraints lets us distinguish those from tracks made by noise. The language of hypothesis testing lets us determine and apply the mathematical form of the constraints given tracks of any length. Given flashes from two and three consecutive frames, the constraints are simply limits on the velocity and acceleration of the track. For longer tracks, the constraints do not have quite so simple an interpretation but turn out to be quite simple to apply.

### Results

Figure 5 demonstrates the effectiveness of the technique on two simple scenes. The image on the left is the true map of what's in the frame, with objects like the road stripped out for simplicity. The image in the middle shows what we'd get from conventional processing with a reasonable threshold, and on the right, what we get after the firefly processing. The firefly process detects objects nearly 300 m away. State-of-the-art LIDAR systems have ranges of about 150 m.

Table 1 shows the detection (%) and number of false positives (per frame) obtained from the firefly process and the conventional processing (MF stands for matched filter). The detection threshold is set so that we have 99.9% confidence based on pre-collected statistics that a particular peak corresponds to an object. However, the detection rate is very low. Using the track constraints helps quite a bit.

Table 1. Numerical Results

	Cluster Size	Length	Detect %	False Pos.
Firefly		3	67.7	3.14
		4	65.6	0.12
		5	61.5	0.04
		6	58.8	0.02
		7	55.1	0.03
		8	51.6	0.02
		9	46.9	0.00
MF Peak 99.9%		10	40.6	0.00
		1	19.1	52.00
		10	18.0	14.34
		20	13.1	0.00
	50	6.7	0.00	0.00

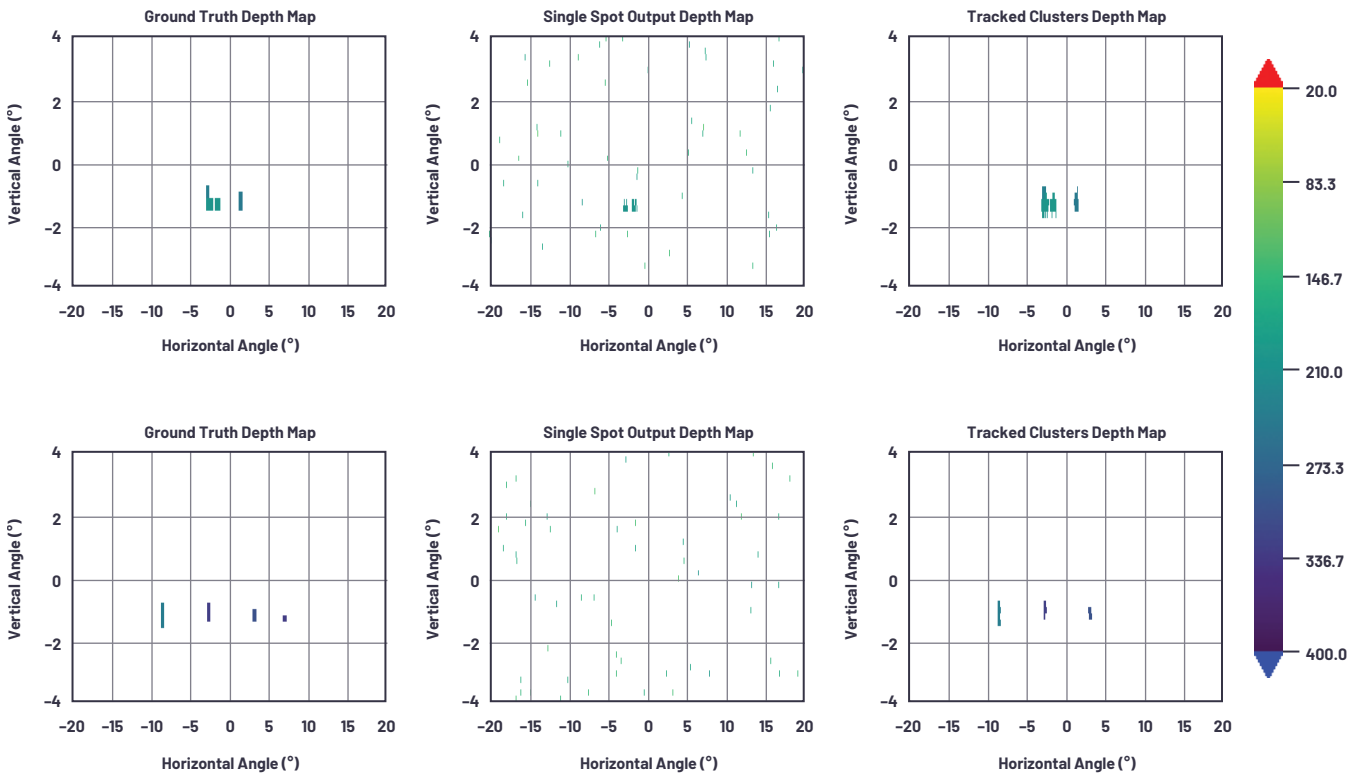


Figure 5. Examples of detecting by tracking using the firefly process.

## Conclusion

The firefly process describes a boundary on how objects can move—that is, it details the constraints, not on the detector or signal chain, but on the object they are measuring. We believe that its power to improve the detection rate holds an important lesson: conventional detection and signal chain problems can often be greatly improved by utilizing constraints and ideas that come from outside of the system we are designing. We hope to continue to leverage such insights in designing smarter and more sophisticated signal chains, and continue to gain advantages from exploiting unconventional constraints whenever we can.

The author would like to thank Jennifer Tang, Sefa Demirtas, Christopher Barber, and Miles Bennett for their contributions to this article.

## Technical Reference

Tang, Jennifer, Atulya Yellepeddi, Sefa Demirtas, and Christopher Barber.

“Tracking to Improve Detection Quality in LIDAR for Autonomous Driving,” IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), May 2020.

## About the Author

Atulya Yellepeddi is a research scientist in the Analog Garage. He earned his Ph.D. in electrical engineering and applied ocean science and engineering from the Massachusetts Institute of Technology and Woods Hole Oceanographic Institution. In his time at ADI he has worked on developing the signal processing and inference chains for a variety of technologies, including time of flight, LIDAR, radar, and audio technology. His research interests include adaptive signal processing, learning by exploiting signal structure, inference, and information. He is the author of several patents and peer-reviewed publications both during and prior to his time at ADI. He can be reached at [atulya.yellepeddi@analog.com](mailto:atulya.yellepeddi@analog.com).

Engage with the ADI technology experts in our online support community. Ask your tough design questions, browse FAQs, or join a conversation.



Visit [ez.analog.com](https://ez.analog.com)