Selecting a LIDAR System

Robert Clark
Acuity Technologies

Laser radar or LIDAR (Light Direction and Ranging) systems use lasers to create distance maps of the environment around them. Different products have different strengths, and knowing the significance of each specification is critical in choosing the right LIDAR for the job. Figures of merit for mapping LIDAR systems include minimum and maximum range, range accuracy, angular accuracy and resolution, laser spot size, field of view, sample rate, angular scan speed, airborne particulate rejection, and eye safety class. These often must be traded off against one another due to basic constraints of physics. This article provides an overview of primary performance specifications and their significance to help users select the best system for their task.

LIDAR may refer to distance mapping or atmospheric analysis systems. In the former, pulses of precisely aimed narrow-beam laser light are transmitted and reflect back from surfaces. The time between pulse transmission and detection of a reflection, measured to tens of picoseconds, gives distance. Since light travels about 1 foot per nanosecond, resolution in the millimeter range is possible. Using beam scanning or other methods, full 3D images can be built up. These are used in a variety of applications from autonomous navigation and mapping to profiling structures and monitoring industrial processes.

Field of View

A LIDAR typically sweeps the beam direction while operating, giving a collection of distance measurements within the LIDAR's field of view (FOV), the range of horizontal and vertical angles through which it can capture data. Moving mirrors may be used to scan in one or two axes, or the entire laser assembly may move, or scanning may be a combination of both. A two-axis scanning LIDAR can capture detailed shape information in the horizontal and vertical directions from a stationary location. One-axis scanners capture distance and angle information along a single scan line, and rely on relative motion of the scanner and the target surface if a full image is to be built up. A hybrid approach is the use of one scanning axis with multiple laser beams, or channels, projected at different angles. This limits the resolution in one axis and relies on motion for full image formation, but can increase sample rate relative to a single-beam system.

The field of view may be anything from a small window in a "forward looking" system to essentially full spherical coverage of 360 by 180 degrees. Typical fields of view are 360 degrees in the horizontal and 30 to 120 degrees in the vertical. Design tradeoffs include size and weight, scan speed, laser power and eye safety, and system complexity.

Maximum Range

The maximum distance that can be measured is generally important to LIDAR users. Factors affecting the maximum range are laser power, target surface diffuse reflectance, and the amount of ambient light coming from the target surface. The reflected laser energy must be sufficient to trigger the pulse detector, which will have some limiting threshold to mask out noise from ambient light. At extreme ranges, this means that some actual range returns may be missed. One measure of maximum range is the point at which the probability P_D of detecting a return is 50%, while the probability P_{FP} of a false positive from noise remains acceptably low. Target reflectance should also be included in a maximum range specification. Since the amount of a reflected laser pulse that is captured at the detector decreases with the square of the distance, a 90% reflective surface might be detectable from three times farther away than a 10% reflective one.

The trigger threshold, whether explicitly set or implicit in signal processing, represents a tradeoff between maximizing the probability of pulse detection and minimizing the probability of a false positive. Different signal detection thresholds may be optimum for different applications and conditions. Some applications may need only highly reliable data and not be concerned with obtaining the maximum sensitivity. Others may require maximum sensitivity but be able to tolerate some spurious signals, perhaps filtering them in software.

Also important is the degradation of accuracy at extreme range. A laser pulse has a short but nonzero rise time and a strong reflection will trigger detection closer to the start of the return pulse than will a weak one, though this can largely be corrected. However, the measured return is always a combination of signal and noise so when the pulse strength is low and close to the noise level the exact time at which the signal crosses the threshold varies somewhat. A shorter laser pulse with a faster rise time will generally result in measurements that are less sensitive to return signal strength and have less variation at long distances. LIDAR pulses are typically from a fraction of a nanosecond to 10 ns long, as measured between the full width half maximum (FWHM) points, the half-peak-power points of the leading and trailing edges.

Minimum Range

When a laser pulse is transmitted, some of the light may reach the receiver after being scattered off the transmit optics or external window. This initial signal must be masked or rejected at some point in the data collection chain in a way that still allows detection of nearby external surfaces. In general, the shorter the pulse the lower the minimum measurement range can be, since the pulses must be sufficiently separated to be distinguishable. In addition, detectors may briefly become "blinded" by a strong pulse. The geometry of the optics may also prevent light reflected from very short distances from reaching the detector. A LIDAR may therefore have a minimum range, and closer targets may give no return or an inaccurate measurement.

Angular Accuracy and Resolution

Angular accuracy is affected by the mechanical rigidity and stability of the system and the accuracy and resolution of angle measurements. High resolution angular encoders can provide very precise angle information, but the effective resolution is also limited by the laser beam diameter and divergence. On moving platforms, accurate image formation depends good information about the platform motion, since LIDAR images take time to acquire. This may be acquired with an inertial measurement unit (IMU) closely coupled to the LIDAR. Angular resolution may be expressed in degrees, microradians, minutes or seconds of arc, or milligons (1 milligon = .0157 milliradian).

Laser diodes used in LIDARs typically emit from a relatively large, non-circular area and have complex beam intensity and divergence profiles, which puts a lower limit on the size of the laser spot on a distant target. This can be improved by expanding the initial laser beam diameter, but the spot size is then larger for short range measurements. If the divergence angle of the laser beam is much larger than the encoder angular resolution, the laser beam divergence will be the limiting factor in the system resolution, and better encoders will be of relatively little value.

Simple transmit optics can also result in a spot on target surfaces that is many times larger in one dimension than in the other. Ror high resolution applications the spot shape, or eccentricity, may be important, and laser-specific optical design can circularize the beam.

Another consideration regarding laser beam divergence is that of immunity to bright ambient light as discussed under **Maximum Range**. Higher levels of ambient light can reduce the sensitivity and maximum range of a LIDAR system. Small detectors or pinholes are also used to pass focused light only from the location of the laser spot. The smaller the laser spot, the smaller the acceptance area at the receiver can be, which reduces ambient light effects.

Higher end LIDAR systems often use other types of lasers such as fiber lasers or solid state lasers. These typically emit circularly symmetric beams that can be collimated to maintain a small diameter over long distances and are also capable of higher per-pulse energy, which allows longer distances to be measured.

Safety

The main hazard from lasers is the risk of retinal damage. Section 21 of the Code of Federal Regulations (CFR) Part 1040 in the US and IEC 60825 in Europe define the safe limits to laser radiation exposure, for times from nanoseconds to hours for pulsed and continuous beams, stationary and scanned, as a function of beam power density, wavelength, and source size. Lower limits on LIDAR scan speeds are often determined by the time limits for safe exposure. For products emitting laser light, manufacturers must file initial and periodic reports with their country's responsible agency. The lowest classification is Class I, which indicates essentially no hazard in normal use. LIDARs used in public areas without safety controls should all fall in this class.

The laser wavelength strongly affects the power level for eye safety. For some time domains, lasers with wavelengths around 1.55 microns can safely emit up to 100 times more power that those emitting wavelengths near or in the visible range. These include some fiber laser and solid state lasers. Although considerably more expensive, these can be used to measure longer ranges.

Range Accuracy

Absolute accuracy in LIDARs has two primary error components. The first is the fundamental distance accuracy limit. A measurement may be affected by factors such as the strength of the signal reflected from a surface, the ambient and system temperature, the amount of background light present, and the quality of the LIDAR's calibration algorithm.

The second factor is the variation, or standard deviation, in group of measurements taken to the same point under the same conditions. This is generally larger when the return signal strength is low and for longer distances and low reflectance surfaces. Causes of variation include noise from ambient light, detector and amplifier noise, laser pulse power fluctuation, and the resolution of the timer used to measure the round-trip pulse time. For a group of measurements with a Gaussian distribution and a standard deviation (1σ) of 5 mm, 68.2% of the measurements will lie within 5 mm of the average of a large set of repeated samples, 95.4% will lie within 2σ, or 10 mm, and 99.8% within 3σ, or 15mm. A $\pm$ error limit is sometimes defined as the $\pm 3\sigma$ point.

If an application requires high accuracy but there is time to acquire and average repeated measurements of the same point, this can reduce the effect of the variation in measurements. Most LIDARs have the ability to trade image acquisition time for better precision by oversampling and averaging. This may be done by making the pulse rate high relative to the scan speed divided by the beam divergence so that successive samples overlap, or the same location may be scanned repeatedly. This reduces the effective sample rate but improves the depth resolution in proportion to the square root of the number of samples averaged. However, the absolute accuracy may still be limited by the first set of considerations mentioned.

Sample Rate

In general the higher the laser pulse repetition rate the better the final result. More tightly spaced points and more range averaging can be used to increase the information acquired in a given time. This is typically important for dynamic applications such a navigation and mapping from a mobile platform or for industrial processes with high throughput. In static applications such as surveying and structure capture the sampling time may be short relative to the setup time and a slower system may be adequate.

Sample rate is typically limited by the laser's power limit, laser eye safety, and the range ambiguity interval. The range ambiguity interval is the speed of light divided by the pulse rate. At measuring distance intervals greater than this multiple pulses may be in flight at once. This can can result in reflections being received out of order, leading to large measurement errors. Thus a system with a 500 foot range window must generally sample somewhat less than 1 million times per second, though various techniques can also be used to mask or disambiguate return signals in sophisticated systems.

Fog and Dust Rejection

Since many LIDARs have much more sensitivity than is needed to detect objects at short range, fog, dust, or other particulates can cause false range returns from pulse rangefinders. This can be ameliorated by reducing the detector sensitivity for the time interval at which short distance returns would occur. As the time and implied distance increases, the sensitivity can be increased. The light blocked by the particulates will reduce the maximum detection range, but a properly defined rejection mechanism will not add to this reduction since sensitivity is back to the maximum level by the time a weaker long distance return signal arrives. Systems may also provide an adjustable version of this sensitivity adjustment that allows increased rejection in dusty or foggy conditions while maximizing detection of hard to detect objects at intermediate ranges in clear conditions.

Multiple Returns from one Pulse

If a pulse falls partly on one surface at the edge of an object and partly on a more distant surface or if particulates or flying debris intercept part of the beam, one of several things may happen. If the surfaces are separated by a large distance, two return pulses may be detected. LIDARs may be designed to record multiple return pulses, only the first pulse, or only the last pulse. For collision avoidance the first pulse may be most important, but for terrain mapping the last may be selected to eliminate partial returns from foliage. If the surfaces are close together relative to the “length” of the pulse (pulse duration multiplied by the speed of light) the return signals will overlap and a distance between the two surfaces will typically be recorded.

Grayscale Intensity Imaging

A LIDAR may provide a grayscale image as well as a range image using the ambient light level and/or the strength of the return pulse at each point. This allows grayscale (“black and white”) images to be displayed or overlaid on the range image. When pulse signal strength is used, the intensity measurements can be normalized for distance.

All-Weather Operation

LIDARs intended for full time all-weather tasks such as autonomous vehicle obstacle avoidance will have a case that is impervious to the environment and will be able to operate over a wide temperature range. Whether the LIDAR will measure effectively in all weather is another matter. A relatively small diameter outgoing laser spot can be significantly deflected by a drop of water on the optical window, and some external protection may be needed to keep it reasonably dry. Miniature windshield wipers have even been spotted on some systems.

Technical Details

The collection and processing of range signals has several stages. First, optical filtering and focusing screens out most ambient light and collects laser returns at a fast optical detector. The detector’s electronic output is then amplified. In some designs, a fast analog to digital converter samples the signal at a high rate (1GHz or higher) and the data stream is analyzed to identify the return signal. This approach is often used with fairly long pulses and provides sophisticated signal processing at the cost of relatively complex hardware and software. In other systems the pulse arrival is timed with high speed electronics such as a TDC (Time to Digital Converter). Time gating may be used to accept only signals from a specific window between minimum and maximum distances. Other information such as the strength of the signal and ambient light level may be collected. In software this data is used to obtain a calibrated distance measurement, and reject any spurious signals.

Wrapup

LIDAR performance has several aspects, and not all applications need maximum performance in all of them. Identifying the technical and economic factors that are most important in a particular situation will help in selecting the right LIDAR system the first time.

Robert Clark is president of Acuity Technologies. He has been involved in rangefinder and LIDAR development since 1990. Mr. Clark can be reached at rclark@acuitylidar.com.