HIGH DEFINITION LIDAR

VELODYNE'S HDL-64E:
A HIGH DEFINITION LIDAR™ SENSOR
FOR 3-D APPLICATIONS

White Paper / October 2007

www.velodynelidar.com

VELODYNE'S HDL-64E: A HIGH DEFINITION LIDAR™ SENSOR FOR 3-D APPLICATIONS

Abstract

Velodyne has developed and produced a High Definition Lidar (HDL) sensor designed to satisfy the demands for autonomous vehicle navigation, stationary and mobile surveying, mapping, industrial use and other applications. The Velodyne HDL unit provides 360-degree azimuth field of view and 26.5-degree elevation field of view, up to 15 Hz frame refresh rate, and a rich point cloud populated at a rate of one million points per second. This White Paper first explores various technologies for 3-D lidar sensing, then discusses the design principals and construction of the Velodyne HDL-64E lidar sensor.

3-D Sensor Technology Overview

The use of light pulses to measure distance is well known. The basic concept is to first emit a light pulse, typically using a laser diode. The light travels until it hits a target, when a portion of the light energy is reflected back towards the emitter. A detector mounted near the emitter detects this return signal, and the time difference between the emitted and detected pulse determines the distance of the target.

When this pulsing distance measurement system is somehow actuated, a multitude of points (called a "point cloud") can be collected If no targets are present, then the light would never return. If the light was pointed downward, then the ground would provide some signal return. If a target was positioned within this point cloud, a notch would be seen in the rendered data. The distance and width of the target could be determined from this notch. When this collection of points is rendered, the point cloud begins to resemble a picture. The denser the point cloud the richer the picture becomes.

There exist a number of commercial products that can capture distance data in a 2-D (i.e. single plane) point cloud manner. These instruments are often used in industrial applications and are often repurposed for surveying, mapping, autonomous navigation, and other uses. Most of these devices rely on the use of a single laser emitter/detector pair combined with some type of moving mirror to effect scanning across at least one plane. This mirror not only reflects the emitted light from the diode, but also reflects the return light to the detector. Use of a rotating mirror in this application is a means to achieving typically 90 – 180 degrees of azimuth view while simplifying both the system design and manufacturability, as there is only one moving part, the mirror.

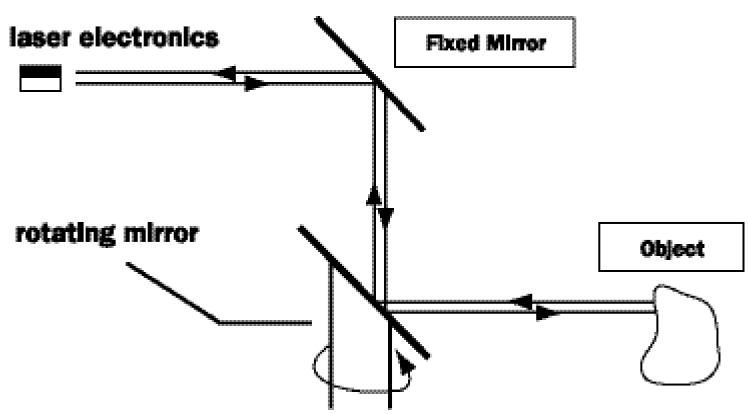

Figure 1. Single emitter/detector pair rotating mirror lidar design.

Such devices are often used in industrial applications, as shown in Figure 2. Note the scan lines emitting from the unit - the spinning mirror allows the single laser emitter/detector assembly to be aimed along this plane via the use of the rotating mirror.

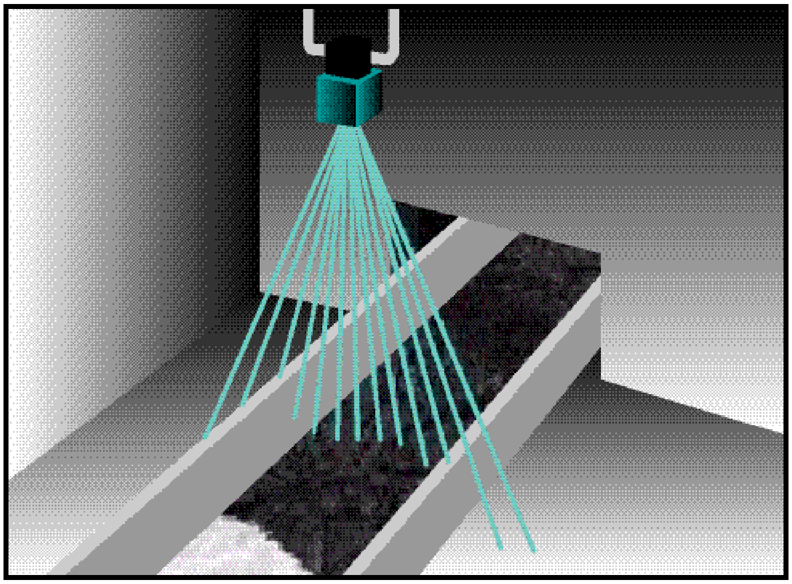

Figure 2. Industrial application of a conventional single emitter/detector pair laser.

Many applications require more data than just a single 2-D plane. The 2-D point cloud must be expanded to form a 3-D point cloud, where multiple 2-D clouds are used, each pointing at a different elevation angle. One method of achieving this third dimension is to move the emitter-detector pair up and down while the mirror rotates (sometimes referred to as "winking" or "nodding"). This will generate elevation data points, but it will reduce the number of azimuth data points, thus the point clouds are less dense, resulting in a lower resolution system.

There also exist "flash lidar" units. These units operate by simultaneously illuminating a large area, and capturing the resultant pixel-distance information on a specialized 2-D focal plane array (FPA). Such sensors are complicated and difficult to manufacture, and as a result not widely deployed commercially. However, it is expected that they may someday replace mechanically actuated sensors, as they are solid state and require no moving parts.

3-D point cloud systems exist in several configurations. However, the needs for autonomous vehicle navigation, mobile surveying, and high speed/high density image capture place unrealistic demands on current systems. For example, there are numerous systems that take excellent pictures, but take several minutes to collect a single image. Such systems are unsuitable for any mobile sensing applications. There are also flash systems that have excellent update rate, but lack field of view and good distance performance. There are single beam systems that can provide useful information, but do not work well with objects that are either too small or fall outside the unit's field of view. In reality, to adapt to the greatest number of uses for a lidar sensor it is necessary to see everywhere around the collection point - a full 360 degrees. In addition, the processed data needs to be presented to the user in real time, so it is necessary to have a minimum of delay between the data gathering and rendering the imaging. For example, in the arena of autonomous navigation, it is generally accepted that human response time is in the several tenths of a second. Therefore, it is realistic to provide the navigation computer with a complete fresh update minimally ten times a second. The vertical field of view needs to extend above the horizon, in case the car enters a dip in the road, and should extend down as close as possible to see the ground in front of the vehicle to accommodate dips in the road surface or steep declines.

Other mobile sensing applications beyond autonomous navigation currently suffer from a lack of sensor data density and frame rates given the vast amounts of area to be captured. These applications include highway surveying and infrastructure assessment, rail surveying and assessment, and environmental image capture (mapping). Currently, stationary sensors are being retrofitted for mobile sensing tasks, or stationary sensing techniques are being used where a mobile sensor is called for. Such approaches represent a higher cost in terms of time, or risk to the surveyor as they accommodate the slow collection rates of the sensor in use, such as blocking off a lane of highway in order to survey the road surface.

HDL-64E Product Description

Velodyne has developed and produced a patented High Definition Lidar sensor, the HDL-64E, shown in Figure 3, designed to satisfy the demands for mobile sensing applications.

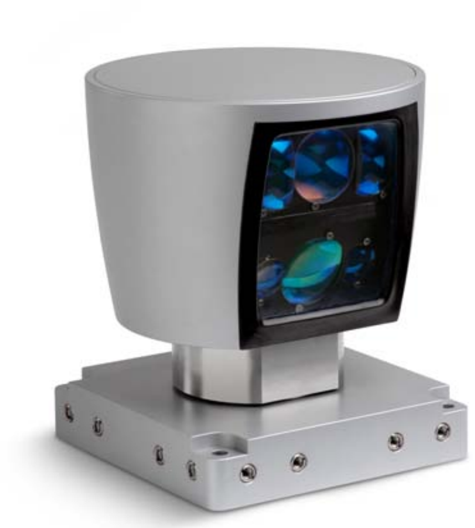

Figure 3. Velodyne HDL-64E.

The HDL-64E is the second generation laser sensor produced by Velodyne. The first generation prototype unit was used in the 2005 DARPA Challenge as part of Team Digital Auto Drive (DAD) where it enabled the Team DAD Toyota Tundra to complete 25 miles of the course before suffering a non-sensor-related mechanical failure. As a result of this success and the capabilities demonstrated by this sensor, including resolution, range and output data rate, the HDL-64E has been purchased and incorporated into over 12 teams' autonomous navigation system designs for the 2007 DARPA Urban Challenge.

The HDL-64E features 64 emitter-detector pairs, each aligned to provide an equally spaced 26.5 degree elevation field of view, spanning from +2 degrees to -24.5 degrees. The entire optical assembly rotates on a base to provide a 360-degree azimuth field of view. The unit has a range of 120 meters and typically has a distance error of less than one inch.

The unit rotates by default at 600 RPM (10 Hz), but can be instructed to run between 300 and 900 RPM on the fly. The rotation rate is controlled by sending a simple text command through a computer serial port interface. This same serial port can also be used to upgrade the firmware stored inside the sensor assembly. The unit outputs a constant one million distance points per second regardless of rotational speed, providing horizontal angular resolution of as little as .05 degree. The unit is both statically and dynamically balanced which minimizes any vibration felt by the user and also provides for a stable rendered image. This performance, in terms of point cloud density, frame rate, and field of view is superior to any other sensor presently available.

The HDL-64E lasers each emit an optical pulse that is five nanoseconds in duration, as measured between the 50% points of the optical peak amplitude, with a maximum peak power output of 60 watts. Each laser is command charged just prior to triggering the laser using a fly back charging technique. Using this charging method, only a low voltage connection is required for the laser assembly. The high voltage necessary to create the roughly 30 amps of peak current discharged through the laser is generated by this fly back circuit located on the laser assembly. The output laser light is focused using a lens. When the light strikes a target, a portion of the light reflects towards the source. This return light passes through a separate lens and a UV sunlight filter. The filter acts to decrease the amount of energy introduced by the sun, which otherwise would tend to raise the noise floor, degrading the signal-to-noise ratio, thereby decreasing the overall sensitivity of the system. This receive lens focuses the return light on an Avalanche Photodiode (APD), which then generates an output signal relative to the strength of the received optical signal. The laser and APD are precisely aligned at the Velodyne factory to provide maximum sensitivity while minimizing signal crosstalk, thus forming the emitter-detector pair.

The system varies the amplitude of the laser pulse based on the return signal strength, as detected by the APD and an amplifier circuit. There are two primary purposes for this automatic power control. First, the system can minimize the laser power output to the necessary level which decreases the heating inside the unit, thereby increasing reliability. Secondly, the power adjustment keeps the detectors out of saturation. When a detector receives too much optical energy, it will saturate, thereby distorting the detected signal. Also, when a detector saturates, it takes a substantial time for the device to recover. This saturation recovery period may exceed the pulse repetition time, thereby distorting the next detected signal. On the other hand, when a signal is near the noise floor, making it difficult to detect, the power level can be raised using this automatic power control. This may be the case when targets are approaching a specified range of 120 meters or when a low reflective target, such as a flat black target, is within the unit's field of view.

The detector output is amplified and passed to a high speed (3 GHz) Analog-to-Digital Converter (ADC), which digitizes in the incoming analog detector signal and then sends this data to a Digital Signal Processor (DSP), which uses a proprietary algorithm to analyze the digitized waveform to determine the time of the signal return. The short optical pulse in addition to the high bandwidth signal processing and algorithm implementation provide for the high resolution of the system.

The emitter-detector pairs are divided into two 32-laser banks. The upper bank (called the upper "block") is physically located near the top of the unit, and is directed at the higher half of the elevation angles. The lower bank is located below the upper optical block, and is intended for the lower half of the elevation field of view, and is therefore looking at shorter distances than the upper block. Because the upper block is directed at higher elevation, and thereby further distances, the angular distance traveled between optical pulses is larger than the lower block.

The sensor outputs data to the user through a standard 100BaseT Ethernet port. Data is continuously streamed out of this port at a frame rate equal to the rotation rate (600 rpm would produce a 10 Hz frame rate), providing over one million distance points per second. Included in these Ethernet data packets are the distance and intensity data for each emitter-detector pair and the angle to which the distance and intensity data pertain. This data can then be captured using a standard Ethernet packet capture utility such as Wireshark, rendered in a computer based program, such as the Velodyne Digital Sensor Recorder (DSR), or for example, can be processed by an autonomous navigation system to develop a cost map that can then be used to identify obstacles, find drivable road, and ultimately determine steering, braking and acceleration.

Figure 4 shows a sample still frame display of the HDL data using DSR. Near the center of this image is a solid white dot, which indicates the position of the sensor. Each point cloud generated by the different emitter-detector pairs is represented by a different color. The summation of these 2-D point clouds creates the 3-D image – for example as the unit spins the collection of points from a single emitter/detector pair over flat ground appears as a continuous circle. This sensor is mounted on the cab of a truck, and the rear bed of the truck can be seen in the image, just below the solid white dot. In front of the truck (and sensor) are two vehicles, a truck which appears to be making a right hand turn, and a car passing on the cross street. Further past this car, the image shows a road divider, consisting of a concrete gutter, ground and trees. To the left and right of the sensor, the image shows a slight undulation in the display, which corresponds to a concrete road gutter and a sidewalk that separates the road from a field. There are also shrubs located around the sidewalk. Data can also be seen to the rear of the sensor, as there is clearly another vehicle behind the truck in this image.

Figure 4 also shows that there are no breaks in the circular data around the truck in any of the point clouds. This is an indication that the laser pulse repetition rate and upper block to lower block duty cycles are configured properly for the sensor. A repetition rate that is too slow would result in each of the circles would appear as dotted lines. It should be noted that one can zoom in far enough to see the individual data points collected by the emitter-detector pairs, but much like viewing a television, if you zoom in close enough to see the pixels, the larger picture can no longer be seen. The only areas of blanking, where there is no data, are between the point clouds or where a shadowing effect occurs, where a target is in the optical transmit path, and thus no information can be obtained from behind the target. The blanking behind the rear bed of the truck is an example of this shadowing effect.

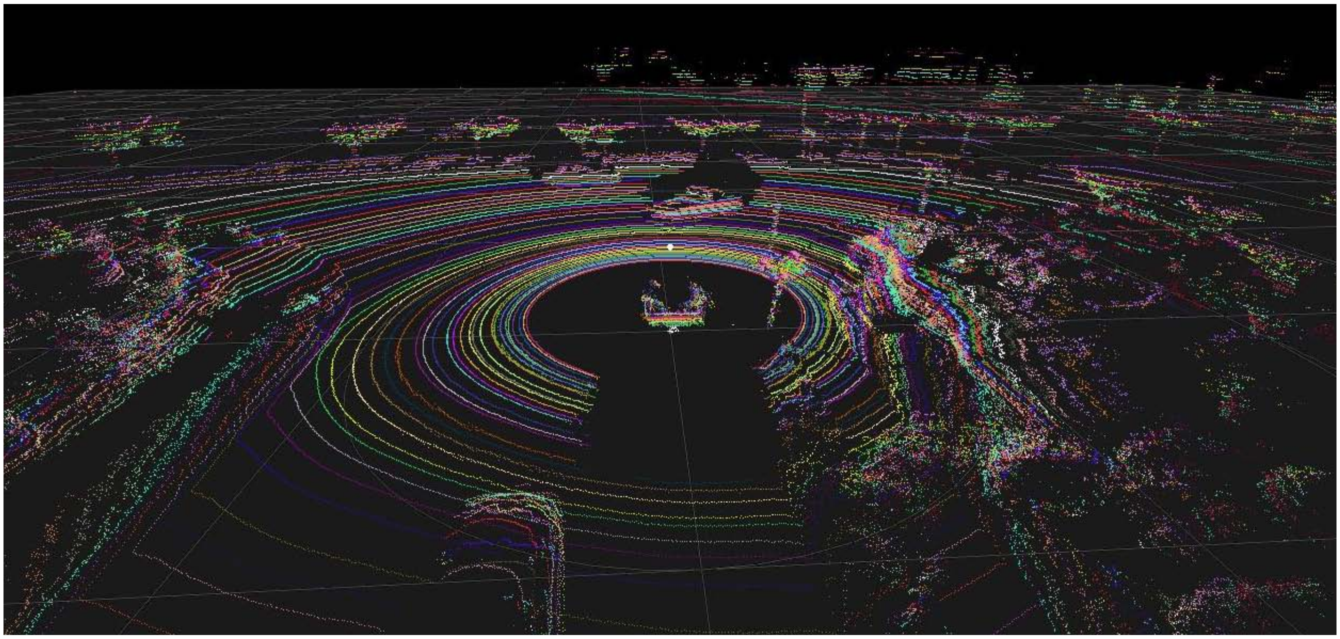

Figure 4. Digital Sensor Recorder Display of HDL-64E Data.

It should also be noted that this unit can be mounted up to 90 degrees off its vertical axis, effectively mounting the unit sideways. The net result of this mounting configuration is to interchange the field of views of the sensor, so that the azimuth would have a 26.5 degree field of view, while the elevation field of view would grow to 360 degrees. This mounting scheme has been used successfully in surveying and mapping type applications.

This sensor is also rated as Class 1M device, meaning it is safe to view with the naked eye, but should not be viewed with any device that has an amplifying viewing optic, such as binoculars. The unit is also waterproof, so when operating this sensor in the rain, no water will penetrate the exterior, preventing potential water damage to the sensors interior assemblies. The unit has also undergone shock and extreme temperature testing with the design goal of making the unit suitable for extended automotive use.

CONCLUSION

Velodyne's HDL-64E is a High Definition Lidar capable of acquiring a large volume of high resolution 3-D data. The HDL-64E features a unique combination of high resolution, broad field of view, high point cloud density and an output data rate superior to any available Lidar sensor in the marketplace today. As such, it is the ideal building block for applications such as autonomous vehicle navigation, infrastructure surveying and mapping display and retrieval, as well as many other applications requiring 3-D data collection.

Velodyne Acoustics, Inc.
345 Digital Drive
Morgan Hill, CA 95037
408.465.2800