A Low-Cost Laser Distance Sensor

Kurt Konolige, Joseph Augenbraun, Nick Donaldson, Charles Fiebig, and Pankaj Shah

Abstract – **Many indoor robotics systems use laser rangefinders as their primary sensor for mapping, localization, and obstacle avoidance. The cost and power of such systems is a major roadblock to the deployment of low-cost, efficient consumer robot platforms for home use. In this paper, we describe a compact, planar laser distance sensor (LDS) that has capabilities comparable to current laser scanners: 3 cm accuracy out to 6 m, 10 Hz acquisition, and 1 degree resolution over a full 360 degree scan. The build cost of this device, using COTS electronics and custom mechanical tooling, is under $30.**

I. INTRODUCTION

One of the most common tasks for mobile robots is to make a map and navigate in an environment. To do so, the robot needs to sense its environment in an efficient way, looking out to some distance to find obstacles and build a map that is useful for performing tasks such as vacuuming or delivery.

While there are many sensors that could be used, laser distance sensors are currently the standard sensor in indoor and outdoor mobile robots. The main reason is the utility of the data: an LDS returns distance to objects in its field of view, unlike (for example) vision sensors, which need complicated and error-prone processing before distances are measured. And unlike other distance sensors such as sonars or IR sensors, an LDS is capable of fine angular and distance resolution, realtime behavior (hundreds or thousands of point measurements per second), and low false positive and negative rates. Efficient algorithms exist for mapping and localization using LDS scans [5][8].

While LDS devices are ubiquitous in research robotics, their high cost has kept them from appearing in consumer robotics such as robot floor cleaners. The Electrolux Trilobite, one of the only cleaners to make a map, relies on sonar sensors [13].[2] The barrier to using LDS technology is the cost. The two most common devices, the SICK LMS 200 [1] and the Hokuyo URG-04LX [1], cost an order of magnitude more than the simplest robot cleaners.

In this paper we describe a compact, low-cost LDS that is as capable as standard LDS devices, yet is being manufactured for a fraction of their cost: the Revo LDS. Figure I-1 shows the prototype Revo with its cover removed. It has the following characteristics:

1. Eye-safe (Class I or II).
2. Works under standard indoor lighting conditions, and some outdoor conditions.
3. Measures a full 360 degree planar scan.
4. Has a range from 0.2m to 6m.
5. High resolution: range error < 3 cm at 6 m, angular resolution of 1 degree.
6. 4000 readings per second (scans up to 10 Hz)
7. Small size, low power (< 2W)
8. Standard, commercially-available components.
9. Low cost: $30 cost to build.

These characteristics make the Revo suitable for consumer products, and open the way for high-performance, low-cost mobile robots. Although all of the Revo technologies have been used in other devices, to date no-one has realized that they could be combined to make a low-cost, high-performance sensor. Achieving the above criteria required innovations in design, algorithms, and integration. The key elements of the Revo are –

- A compact, rigid point-beam triangulation module incorporating laser, imager, and electronics. With a low-cost CMOS imager and a DSP for subpixel interpolation, we get good range resolution out to 6 m

[2] The Samsung VC-RP30W apparently makes a map, but there are no details available on how it does this.

Figure I-1   Revo LDS. Approximate width is 10cm. Round carrier spins, holds optical module with laser dot module, imager, and lens.

with a 5 cm baseline, at a 4 KHz rate. The key insight to the Revo is that high precision is possible with a small baseline, because of the digital image sensor.

- Module rotation to achieve a 360 FOV. Rather than using mirrors to manipulate the beam, the Revo *revolves* the optical assembly to point the beam around a full circle.

No other current device satisfies our requirements. We briefly review several relevant competing LDS technologies that use triangulation.

Structured line devices. These devices use a light stripe laser and offset camera to determine range to a set of points. Because the laser energy is spread over a line, it is difficult to achieve accurate range, especially in the presence of ambient light, or with darker objects. For example, the system of [7] has a maximum 3m range with a 105 degree FOV, and needs a 50 cm baseline to achieve good range precision.

Point scan devices. Many devices exist for 3D scanning of small objects in fine detail ([3], Section 2.1). These typically use a scanning mirror to direct a point laser beam and redirect the laser return to an optical receiver. Such devices cannot be miniaturized, and their cost and mechanical fragility will remain high.

Point modules. Like the Revo device, there exist single-point range modules that could potentially be rotated to achieve a large FOV. Typical are the Sharp IR sensors [9], which use a position-sensitive device (PSD). These devices measure the centroid of all light impinging on their surface. Although modulation techniques can be used to offset some of the effects of ambient light, PSDs do not perform well unless the laser spot has a very strong reflection, limiting their use to ranges of a meter or less [10].

In the rest of this paper, we discuss aspects of the Revo sensor: the modularization of the point sensor, the scanning mechanism, and synchronization. .

II. SINGLE-POINT DISTANCE MODULE

The Revo relies on an innovative laser point sensor module that works on the triangulation principle, using a laser point beam and a digital image sensor, separated by a small baseline. The module incorporates laser, sensor, optics, and computation in a small, rigid package (Figure II-5). It is slightly larger than current IR distance sensors (e.g., the Sharp IR devices [9]), but has much better accuracy and speed.

A. Triangulation Technology

All single-point scanning sensors, such as the SICK and Hokuyo devices, use mirrors to scan the point sensor. These devices are *time-of-flight* distance sensors: they measure the time it takes for light to travel to an object

and be reflected. An alternative technology is *triangulation*: distance to an object is measured by the angle of the reflected light. Figure II-1 shows the basic geometry of triangulation. A laser produces a small point of light, which reflects off an object and onto the image plane of the camera. An ideal pinhole camera is oriented so that the laser beam is parallel to the ray through the center of focus to the edge of the image. This gives a distance measurement from infinity (at one edge of the image) to the distance q_{min} (at the other edge). From similar triangles, the perpendicular distance to the object is

$$q = \frac{fs}{x} \ . \tag{1}$$

The distance along the laser ray depends also on the angle of the laser with respect to the image axis:

$$d = q / \sin(\beta) \ . \tag{2}$$

These equations show the hyperbolic relationship between image distance and object distance that is a property of triangulation. This nonlinear relationship poses problems for determining longer distances – the range sensitivity dq/dx grows quadratically with distance.

$$\frac{dq}{dx} = -\frac{q^2}{fs} \ . \tag{3}$$

For example, if a 1-pixel image displacement corresponds to a 1 cm distance displacement at 1m, then it corresponds to a 4 cm displacement at 2 m.

B. Single-Point Module Design Criteria

The criteria for minimum distance (from Eq. 1) and range resolution (Eq. 3) pull in opposite directions: a small fs product gives a small q_{min}, a large fs has good range resolution.


Figure II-1 Triangulation geometry from similar triangles. The distance to the object is given by the angle of the laser spot in the image. The distance x in the image is measured between the ray parallel to the laser beam and the ray from the object.

3003

The relative weight of *fs* is determined by the image sensor, so we first decide on it. The image sensor should have a short exposure time to improve ambient light rejection (Section II.E), and a large number of pixels for resolution of *x*. We chose a global-shutter CMOS sensor with 752 pixels of resolution and a minimum shutter time of 35μs. Each pixel is 6μm, and we expect to be able to resolve the laser dot to within 0.1 pixel or better.

With these parameters, we can plot the effect of *fs* on range resolution and min distance (Figure II-2). If the min distance is to be 20cm or less, *fs* should be 900 or less. If the range resolution is to be 30mm or less at 6m, the *fs* product should be greater than 700. We pick 800 as the "sweet spot" for the device.

The product *fs* = 800 can be achieved in different ways, but the bias is towards a compact baseline, while keeping the focal length reasonable (larger focal lengths demand longer lenses). With a baseline of 50mm, the focal length is 16mm, and we chose this combination.

Finally we can determine the angle β of the laser relative to the optical axis:

$$\beta = \arctan(f/(376*6\,\mu m)) \approx 82^o \qquad (4)$$

C. Module Calibration

The total error of the device is a function of the device parameters, the error in dot resolution, and the calibration of the device. *Calibration* here refers to all the misalignments that could affect the ideal operation of the device. Because we are using low-cost optical components, the design must account for major inaccuracies. The main ones are laser pointing angle, lens pointing angle, and lens distortion.

- *Laser pointing angle*. The laser must point vertically in a plane parallel to the base of the device, and point horizontally at an 8 degree angle towards the principal ray of the camera. Low-cost laser modules have

typical pointing accuracies of at best 6 degrees. We compensate for pointing angle mechanically, using laser module rotation and a vertical rotation mount.
- *Lens pointing angle*. The diagram of Figure II-1 shows the laser beam and the lens principal ray in the same plane. Generally this will not be the case. Instead, in calibration we search for the horizontal scanline that best corresponds to the laser beam at all distances. We then use 5 scanlines above and below this central line. If the imager is not rotated too much relative to the plane of the laser and focal point, these lines are sufficient to approximate the ideal planar geometry.
- *Lens distortion*. For a low-cost 16mm lens, the distortion will be at least a few percent at the edge of field, even when optimizing for a single wavelength. This is enough to be the major error in distant readings, and must be compensated.

We use a two-step calibration process to deal with lens distortion. First, we fit a 1/x curve using the following procedure.
- Localize the laser dot image to subpixel accuracy (Section II.D).
- For a set of readings at known distances, fit the ideal curve of Eq. 1, weighting distant readings more heavily. There are two parameters to be optimized: the product *fs*, and the pixel offset for calculating *x*.

The 1/x fit yields the curve of Figure II-3. While the data points seem to lie well on the curve, the steep slope at larger distances means that small deviations result in large calibration errors (Equation 3). Figure II-4 shows the errors due only to the 1/x curve fit, at different distances. Both the sensitivity of distance to pixel errors, and the amount of pixel error from distortion, increase at larger distances (towards the edge of the imager). Note that calibration errors at distances less than 1m are less than 2mm.


Figure II-2 Min distance and Range resolution relative to *fs*. The design criteria is to keep min distance less than 20cm, and range resolution less than 40mm. The vertical line is a sweet spot.

Figure II-3 1/x calibration curve. The raw image value is the pixel value of the centroid of the laser dot, interpolated to 1/32 pixel.

3004

Figure II-4 Calibration error as a function of distance. Errors at larger distances are magnified by sensitivity.

To reduce calibration errors, we use a table of offsets for the errors at discrete value of 1m and above, and interpolate between these values for an actual reading. This technique effectively eliminates calibration errors at larger distances, relative to pixel localization noise (see Section II.F).

A calibration is only useful if the device will keep the calibration under conditions of thermal stress and mechanical shock. The physical linkage between lens elements, imager, laser, and laser optics must be rigid and have low thermal distortion. Any relative movement of the chassis that causes the laser dot to deviate more than a fraction of a micron can cause large distance errors, especially at larger distances. The prime culprit is a relative rotation of the laser and imager, caused by thermal expansion or mechanical shifting of the chassis elements linking them.

Figure II-5 shows the frame design for the LDS single-point module. The chassis is machined aluminum, and the lens elements are glass, and focus adjustment for the


Figure II-5 Rigid frame for the LDS single-point sensor module. Imager and DSP mount on circuit board behind chassis. Overall length is 7cm.

lenses is made by screwing it in or out. Both lenses can be locked down by set screws and glue.

D. Laser dot localization

To reduce errors at larger distances, the image of the laser dot must be localized to subpixel precision. We use a simple centroid algorithm for localization. First, the rows in 10-pixel horizontal band are summed. The resultant line image is then differentiated and smoothed, and the center of the dot is found using the maximum value. Finally, the centroid is calculated as

$$\sum_i I(i) \cdot i \Big/ \sum_i I(i) \quad . \tag{5}$$

Table II-1 gives typical results for the centroid method. At short distances, the image of the dot is tens of pixels wide, and large STDs are tolerated. At longer distances, the dot image becomes only a few pixels wide, and sampling effects can be important. The anomalous readings at 2 m and 4 m are probably a result of the dot being near a pixel boundary, a well-known effect in finding the center of a dot [4]. We plan to investigate matched filter methods for better subpixel localization [8][12]. Even with the centroid method, localization is at 0.2 pixels or better for longer distances.

Dist	.25	.30	.40	.50	1.0	2.0	3.0	4.0	5.8
STD	2.3	2.2	1.0	.34	.02	.10	0.05	0.20	0.05

Table II-1 Standard deviation of dot localization in pixels at different distances (m).

E. Ambient light rejection

In most environments, the image of the laser dot is corrupted by ambient light. Two techniques for rejecting this interference are temporal and wavelength filtering [7].

We chose a visible red wavelength (650 nm) for the laser, because it yields slightly higher laser output for eye-safe use, has better imager response, and is easier to debug and calibrate than IR wavelengths. A 20 nm bandpass filter reduces the ambient light flux.

Temporal filtering uses the imager's global electronic shutter to expose the imager just when the laser is pulsing. By using short pulses, the laser power can be increased while still keeping it eye-safe. The ANSI Z136.1-2000 standard [2] for eye-safe lasers does not allow as much overall energy in shorter pulses, so there is a tradeoff between ambient light rejection (favoring short pulses) and imager response (favoring long pulses and higher total energy). Figure II-6 plots the Maximum Permissible Exposure (MPE, or total energy) of each pulse, and the pulse power, against the pulse width.

3005

At the minimum exposure of 35μs, the pulse power can be over 5 mW, which helps to overcome ambient light interference. At longer pulse times, the pulse power drops, while the energy available for reflection grows substantially, and enables darker objects to be seen. The vertical line shows these values at 60μs, our chosen pulse width.

The Revo will work outdoors, even in direct sunlight. The power in sunlight at sea level and 650 nm is about 1.3×10^{-3} mW/mm^2/nm, so a 20 nm filter gives .026 mW/mm^2. The diameter of the laser dot is approximately 3 mm, so the power from the dot is about 0.17 mW/mm^2, a factor 7 greater than direct sunlight. The presence of sunlight does not affect the STD of the dot localization. However, the Revo will not work when the imager stares directly into a light bulb or the sun.

F. Electronics

Processing the images to provide distance readings involves several steps:

1. Pulsing the laser and exposing the imager.
2. Reading out the imager rows.
3. Processing the rows to determine the laser dot centroid.
4. Calculating the distance corresponding to the image centroid.
5. Formatting and communicating the distance measurement.

The block diagram of Figure II-7 shows the main electronic components. The CMOS imager has integrated timing and control, and requires only a frame pulse to start exposure and subsequent readout of 10 rows; the same pulse starts the laser output. The processor, a DSP, streams the image data directly into internal memory, where it is processed to find the dot centroid and map the centroid position to distance. The only external memory

Figure II-7 Block diagram of the main electronic components.

is a small serial flash to hold the DSP programs and calibration data.

All the main components fit on a small PC board attached to the lens module (see Figure II-5). The module uses less than 1W of power in normal operation. Exposure and readout occur sequentially, while processing is performed in parallel with these operations. The primary limitation on speed is the time taken to read out 10 lines. With on-imager binning of lines, it is possible to perform an expose-process-readout cycle in under 0.25 ms, for a read rate of 4000 distance points per second.

G. Performance

We tested the LDS single-point module in two ways:

1. Error vs. distance for a newly-calibrated module, using white targets (>90% reflectance).
2. Error vs. distance for 10% reflectance.

Figure II-8 shows laser dot localization errors for targets with 10% and 90% reflectance. These errors are random errors that arise from trying to localize the position of the laser dot on the imager to sub-pixel precision. The pixel errors are converted to distance errors using the fitted 1/x curve.
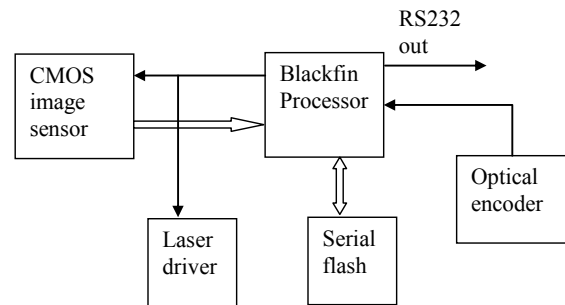

Figure II-6 Maximum pulse power and Maximum Permissible Exposure as a function of pulse width. The chosen width is shown by the vertical line.

Figure II-8 Distance errors for 10% and 90% reflectance.

3006

The 90% reflectance errors are below 3 mm out to 3 m. At 3 m, the sensitivity of triangulation is 78 mm/pixel, so the localization error is less than 0.05 pixels. At increasing distances, the sensitivity starts to create higher errors. The errors at 4 m are an anomaly, caused by our naïve centroid algorithm.

An interesting phenomenon is that the errors do not go down very much below 0.5 m. This is because the apparent size of the laser dot grows, and more pixels become saturated at cloaser distances. Thus, it is more difficult to localize the dot accurately.

Results for 10% reflectance (almost black) were also generated. The errors are significantly higher at larger distances, showing increased uncertainty in localization.

For comparison, the calibration errors to the 1/x curve are also shown. At 1 m and above, compensating for lens distortion is very important in getting good performance. Even with a small baseline, the Revo module exhibits low-error performance out to 6m.

III. SCANNING AND SYNCHRONIZATION

To increase the field-of-view of a single-point distance sensor, it must be scanned. The typical scanning configuration for triangulation sensors users mirrors to deflect the laser beam and return reflections to the image sensor. Such an arrangement is inherently bulky and difficult to calibrate, requiring precise positioning between mirrors, imager, and laser. It is also difficult to achieve full scanning coverage – typically coverage is 180 deg or less.

By contrast, the Revo module is small and rigid enough to be mechanically scanned. In the current device, the module is rotated in a plane, generating a full planar scan at up to 10 Hz. This unique mechanical arrangement, without costly mirrors and consequent alignment problems, enables the Revo to function reliably, while keeping manufacturing costs low. Other arrangements of the module are also possible, e.g., a full 3D scan could be generated by having the module measure not just a single point, but a set of points, or a laser line.

A. Rotational Scanning

The Revo module is mounted on a bearing and spun around an axis midway between the laser and the imager (see Figure I-1). As the module is rotated, the laser is pulsed and a reading is taken at 1 degree resolution. At 10 Hz rotation rate, this is 3600 readings/second, below the maximum rate of 4000 readings/sec.

Power for the laser module is supplied through a 2-wire slip ring on the rotation center. Communication to the module is via a short-range radio frequency modem, at 115 Kbaud, sufficient to send 2-byte data for each reading.


Figure III-1 Rotational geometry of the LDS, with coordinate system centered on the center of rotation *c*.

The geometry of the rotational scan is illustrated in Figure III-1. The laser is offset from the center of rotation; the x,y position of a scanned object, relative to a fixed coordinate system with origin on the center of rotation, is given by

$$x'\,y' = r\cos\varphi, r\sin\varphi$$
$$\alpha = \pi - \beta + \varphi \qquad\qquad . \qquad (6)$$
$$x, y = x'+b\cos\alpha, y'+b\sin\alpha$$

Here the distance *b* is measured along a line through the rotation center and parallel to the image plane (along the segment *s* in Figure III-1).

B. Angular Synchronization

The Revo incorporates a low-resolution optical encoder on the rotating module. A fixed radial black-and-white pattern is read by two reflective sensors on the module. One of the sensors reads an index mark to give the nominal heading of the Revo, while the other reads a 30 cpr pattern for timing the exposure of the laser and imager. Using this technique, the angular displacement of the laser readings is relatively insensitive to variations in motor speed, allowing for cheaper motors and relaxed motor control.

Figure IV-1 shows two overlayed range scans from the Revo. The LDS is in the lower left corner of a small square room, with a few objects (garbage can, table). Note the straight walls that are clearly seen, even out to 3m from the LDS, and at a grazing angle. The readings from the second scan overlay the readings of the first scan almost perfectly. The Revo has been used successfully in mapping and localization experiments with a small cleaning robot.

C. Durability

Mechanical scanning of the optics module raises issues of durability. The center slip is rated at least 1000 hours, good for 3 years' of use at 1 hour/day. The current gear-based outer drive has been tested in continuous use over

Figure IV-1 Two superposed room range scans from the Revo LDS. There are 360 readings at 1 degree intervals. The robot is positioned in the lower left corner; the maximum distance to the far wall is 3.3 m.

several weeks, but because of excessive noise, we are re-designing the drive to have a center bearing. Lifetime of the drive should be determined by the motor lifetime.

IV. CONCLUSION

There are many challenges in transitioning from a proof-of-concept sensor to a consumer product. The cost of the final device is of overriding concern, and compromises must be made in materials and electronics. The key enablers for the Revo are the high rigidity of the laser-to-imager interface and the rapid sub-pixel localization of the laser dot, all using standard low-cost optics and electronics. The Revo should be in volume prototype production by April 2008.

ACKNOWLEDGMENT

REFERENCES

[1] Alwan, M., Wagner, M., Wasson, G., & Sheth, P. Characterization of Infrared Range-Finder PBS-03JN for 2-D Mapping. ICRA 2005.
[2] American National Standard for the Safe Use of Lasers, Z136.l-200, The Laser Institute of America, 1993.
[3] Blais, F. Review of 20 Years of Range Sensor Development. Journal of Electronic Imaging, (13)1 (2004).
[4] Fisher, R. B. and D. K. Naidu. A Comparison of Algorithms for Subpixel Peak Detection. Springer-Verlag, Heidelberg, 1996.
[5] Gutmann, J. S. and K. Konolige. Incremental Mapping of Large Cyclic Environments. In *CIRA 99*, Monterey, California, 1999.
[6] Hokuyo Automation. Scanning laser range finder for robotics.http://www.hokuyo-ut.jp, 2005.
[7] Mertz, C., J. Kozar, J.R. Miller, and C. Thorpe. Eye-safe Laser Line Striper for Outside Use. Intelligent Vehicle Symposium, 2002.
[8] Montemerlo, M. and S. Thrun. Large-scale robotic 3-d mapping of urban structures. In *ISER*, Singapore, 2004.
[9] Sharp Microelectronics. GP2Y0D02YK0F IR Distance Sensor. Datasheet at www.sharpsma.com.
[10] Strobl, K. H. et al. The DLR Multisensory Hand-Guided Device: the Laser Stripe Profiler. ICRA 2004.
[11] Wang, X., J. Gao and L. Wang. A Survey of Subpixel Object Localization for Image Measurement. Proc. ICIA, 2004.
[12] Welch, S. S. Effects of window size and shape on accuracy of subpixel centroid estimation of target images. NASA Technical Paper 3331, 1993.
[13] Zunino, G. Simultaneous localization and mapping for navigation in realistic environments. Lic. Thesis, KTH, Stockholm 2002.