



Last month, this column took a look at how to use the Taos TSL3301 linear image sensor. Adding vision capability to your robot can be exciting. This month, we'll cover a way that your robot can gather useful information about its environment using laser range finding.

This may sound extremely high tech and difficult to do but, in reality, it isn't so hard. There are only a few choices out there when you want to figure out how far you are from objects in your environment.

The old standard is to use sonar. You will often see university robots that have a big ring of gold disks around their perimeter. These are sonar transducers. Sonar is fairly reliable but can be tricked by soft surfaces and tends to round out details such as corners.

If you want to measure short distances, you can use the Sharp GP2D12 infrared range finders. They are good for distances up to 30 inches. Various things can trick them too such as too much ambient light. Another solution is to do stereoscopic range finding. This can be somewhat processor-intensive.

While laser range finding has its faults as well, it seems to be a blend of the best things of these types of sensors. Laser range finders can detect long distances and are able to detect crisp details such as the corners of rooms or possibly the biggest obstacle to mobile robots: chair and table legs.

How it Works

Let's take a look at how laser range finding works. Take a look at Figure 1

for a graphical look at how things are arranged. A laser and an imaging sensor are aligned so that the laser and camera are aimed in the same direction. After a certain distance, the camera will be able to see the laser dot. The farther from the range finder an object is, the higher the laser dot will be on the image sensor. It is really no more difficult than that.

In Figure 1, the gray lines represent the distance detected by each of the 102 pixels of the TSL3301. You can see that the precision that the range finder can resolve distances at is much higher if the object is close to the sensor. The farther that your object gets from your sensor, the less accurately your sensor will be able to resolve its position. This doesn't leave you with much accuracy

using a sensor that only has 102 pixels. Fortunately for you, later in this column you will be shown how to greatly increase the precision that you can resolve longer distances at.

Let's look at what else you can do to increase the accuracy of your range finder. Figure 2 shows what happens if the camera is tilted downward. You might think that this would increase the accuracy of the sensor, but it doesn't. The only difference that this makes is to allow you to detect distances that are nearer to the range finder.


Increasing Accuracy

There are three things that you can do to increase the accuracy of your range finder at longer distances. The first

Figure 1. A diagram of how a laser range finder works.



Figure 2. The range finder with the camera tilted downward.






Figure 3. The range finder with a narrower field-of-view.




Figure 4. The range finder with more distance between the laser and image sensor.

is to use a sensor with more pixels. There is a linear relationship between how many pixels you have and how many distinct distances you can measure. Another way to increase your accuracy at distances is to decrease the field-of-view. Figure 3 shows the result of that. Of course, by doing that, you lose a lot of nearer distances that you could detect before. You can compensate by tilting the image sensor downward, though.

The last thing that you can do to increase your accuracy is to increase the distance between your image sensor and the laser. Figure 4 shows the result of doing that. Once again, tilting

your image sensor can compensate for the loss of nearer measurements. There is a danger to increasing the distance between your image sensor and the laser that you should be aware of. When you make this increase, you also increase the likelihood that something closer to the range finder will obscure its view of the laser dot. This might seem like a non-issue, but it might be good to restrict the distance between the two to three or four inches.

Rolling Your Own Range Finder

In a nutshell, that is all there is to making a laser range finder. Figures 1 through 4 were generated with a piece of software written for this column. You can download it from *SERVO*'s website (www.servomagazine.com) to help you figure out what the best setup for your application would be. Let's look at how you can actually make your own range finder. This design doesn't allow for tilting of the image sensor but works well, just the same. A piece of soft plastic was cut long enough to allow for the desired sensor-to-laser spacing. Measure this distance on the plastic and mark it. The lens hole was first drilled with a hole size that was too small for the lens to pass through, even by forcing it. Next, a second hole was drilled in

the same location that was also smaller than the lens's diameter but big enough to allow you to force the lens into the plastic. Drill this hole only as deep as the thickness of the lens. This hole creates a 'force fit,' which will hold the lens in without the need for any glue.


You should experiment with the proper hole size for the force fit on another piece of plastic. Don't use anything that could damage the lens to force it in there. With the proper hole size, you should be able to insert it with your finger. Last month's column used a special lens setup that allowed you to focus, but this month, we are going to do away with that and figure out the proper distance between the lens and the sensor to get a good focus at most distances.

The laser used to make the test setup for this column was from a cheap laser pen that was found at a drug store. It was easy to mount because — like the lens — it had a circular profile. A hole was drilled that was just slightly too small to fit the end of the pen into. This allowed for a force fit of the laser, as well. Now we will figure out the proper focus distance.

This can be done by holding some tracing paper or other thin paper behind the lens in your range finder. Find the ideal distance and then search around for something that is the same thickness. Make sure to allow for the bit of clear plastic that covers the actual sensing silicon in the sensor chip.

In the test setup, another piece of plastic that was lying around was used. A slit was cut into this piece of plastic to allow light from the lens to hit the sensor but which blocks other ambient light. Before this was glued to the main body of the range finder, the sensor and small piece of plastic were placed in front of the lens and the focus was verified using the image viewing program that can be found on *SERVO*'s website.

Once everything is good to go, glue the plastic with the slit in it to the body of the range finder. You will now glue the image sensor on top of the plastic with the slit in it. This requires a bit of accuracy, so be careful. You want to make sure that the line of pixels points directly towards the laser. If your sensor is rotated, you probably won't be able



to measure all distances. You will also want to make sure that the center pixel is lined up with the center of the lens.

Once you have the image sensor glued into place, you are almost there. Run the image-receiving program to view the results of your handiwork. If you are lucky, the laser will be visible at all distances. It is likely, though, that you won't see the laser at longer distances. You will need to rotate the laser from side to side so that it can be seen at farther distances. It is likely that this won't be very much of a rotation. Once you have it perfect, put some glue onto the laser to hold it in place and let it dry before proceeding.

Ready, Set, Go

Okay, your sensor is now complete. Let's look at what it will take to get it going. The first thing that you will need to do is write some software to find the laser dots. This isn't as hard as you might think. Lasers are very bright so they make a very nice peak even if you turn the gain and exposure time way down. This makes everything else in the room very dark in comparison. This strategy works for most indoor situations quite nicely. Of course, it can't deal very well with extremely bright areas.

A strategy that you can employ in that sort of situation is to take two images in rapid succession. The first image would be with the laser turned on and the second image would be with the laser turned off. Subtract the second image from the first and the only thing left except for the internal noise of the sensor will be the laser dot.

It is very easy for a person to be able to tell where the laser dot is in Figure 6 but how can a computer figure out where the laser dot is? The strategy of simply looking for the brightest area is a pretty good one. If you look at Figure 6, you can see that the laser dot peaks for two pixels. If you took the first pixel that had the maximum value, then you would be fairly close to the actual measurement. Let's look at Figure 7 now. In this case, the distance being measured is short and the laser dot fills up more of the field-of-view. Also, since the focus is adjusted for




Figure 6. A brightness graph and returned image of a laser dot.




Figure 7. The laser dot for a short distance.

distance measuring, the dot is a bit blurry. This results in a large peak.

Obviously, taking the first pixel that peaks first will give you a very inaccurate distance measurement. You can get close by using the middle pixel that has the peak value as your distance measurement.

Using the middle brightest pixel value can get you close to an accurate distance measurement but it doesn't get you all of the way there. In Figure 8, you can see that at longer distances, you don't have multiple pixels that peak. You don't even have a pixel that goes to full brightness. Out at these distances you are going to have terrible accuracy by just using the pixel that has the highest peak. This gives you a clear idea of where the laser is, but at longer distances, the accuracy falls off because of the longer distances between where a pixel's field-of-view intersects the laser dot. Let's look at a way that you can achieve sub-pixel resolution for where the laser dot really is so that you can make much more precise measurements at longer distances.

Figure 9. Pixels that are added to the 'total' variable are highlighted in yellow.




The Brightest Pixel of Them All

This strategy starts out like what was described above where you search for the brightest pixel value. If multiple pixels share the brightest value, then the first one is noted. Now that you have figured out where the first brightest pixel is located, you will create a variable called 'total' and load it with the brightness value of the pixel before the first peak pixel. Now add the peak pixel's brightness value to total.

For each successive pixel after the first peak pixel that has the same brightness value, add its value into total, as well. After the last peak pixel is added, add in the brightness value of the next pixel after the peak. Figure 9 shows the pixels that would be added for a narrow and wide peak.

Figure 8. The laser dot for a long distance.



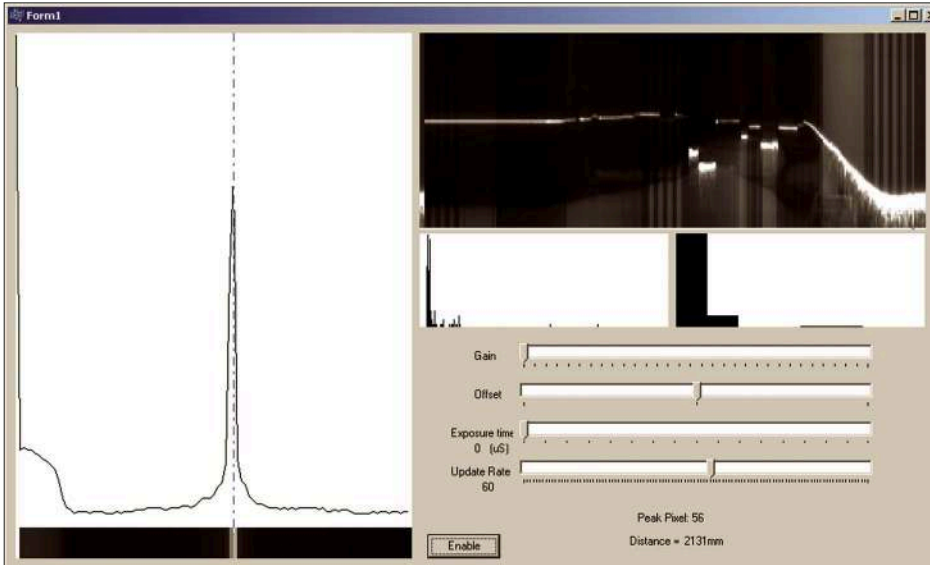


Figure 10. The updated image-receiving program — the vertical bands in the image are due to the flickering of a fluorescent light.

The total variable now contains the total of the values of the peak pixels and one pixel on either side of the peak. Divide this value by two. Now go through the pixels starting at the pixel before the peak and subtract its brightness value from the total variable. If the subtraction results in a negative number, then the center pixel is that pixel. Add that pixel value back into total. This calculation usually gets us the same pixel as what was described before. We can take that value and run it through a lookup table to get an approximate distance. Store that value

in a variable called 'theDistance.'

Here is where this code differs from what was shown before: We will now find the difference in distances between this pixel and the next and store that in a variable called 'pixDif.' Multiply pixDif by the value in total. Now divide by the brightness value of the current pixel. Add this value to what is stored in theDistance. This will be a more accurate estimate of the distance than if you had simply looked up the value and had stopped there. In short, this chunk of code finds the cen-

ter peak pixel and then looks at the pixels on either side of the peak to decide how much to adjust the center location based on their brightness values.

Go the Distance

The last thing that you will need is the actual lookup table to figure out your distances. You can get a pretty good approximation with the software on *SERVO's* website. The code in the lookup table generator gives approximately the correct distance. The range finder built for this column was within 20 millimeters at medium distances using a table generated with this program. Factors such as misaligned sensors or lens distortion will reduce the accuracy of the lookup table.

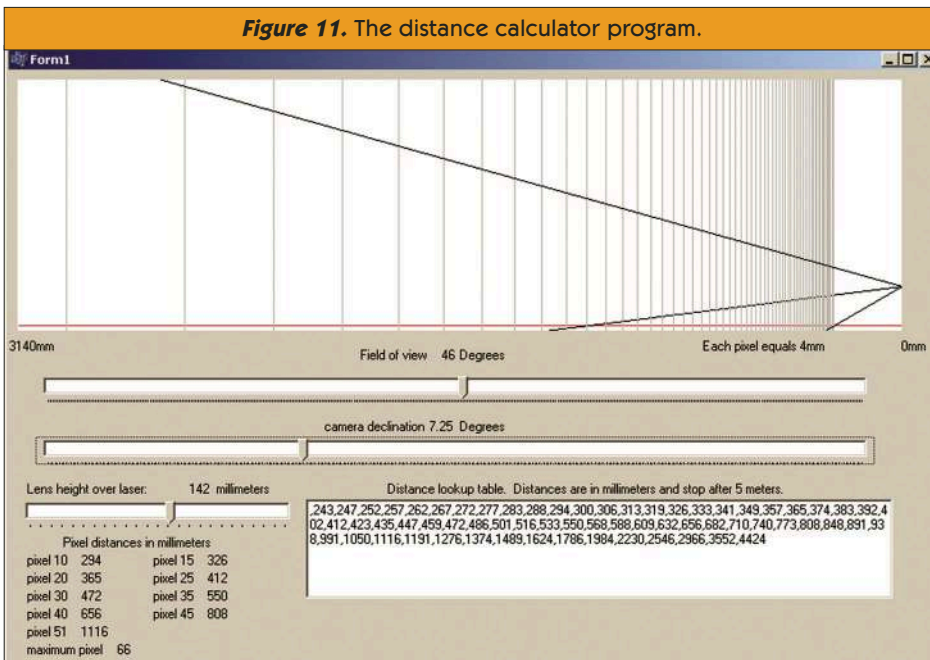
To find the best lookup table using the lookup table generator, you should first use the newer version of the image-receiving program presented last month. This program allows you to do other things such as create 2D pictures. It also shows a 256 and eight bin histogram of the image that it is currently processing. The important part right now is that it tells you which pixel has the peak value.

Clamp your range finder in a vise or otherwise position it so that you can adjust the distance to a target. Now run the image receiving program. Move your target until it has a peak pixel that is far away and is a multiple of five. Measure this distance in millimeters. If you only have a measuring tape that is in inches, then multiply the number of inches by 25.4 to get millimeters. Now move your target until your peak pixel is 10. Measure this distance. It wouldn't hurt to measure a couple other distances for pixels that are multiples of five, as well.

Now switch to the other program that draws a representation of the pixel distances. Set the lens height over the laser and the approximate field-of-view for your lens. On the bottom left of the screen is a list of pixel values and their distances. Play around with the camera declination and field-of-view until your largest distance and the distance to the tenth pixel match as closely as possible with what the program says. Put greater emphasis on getting the longer distance to be correct.

When you have a good match, copy

Figure 11. The distance calculator program.



the lookup table from the bottom right side of the screen and use it in your program. If you happen to have Borland C++ Builder, you could copy this lookup table into the source code of the image-receiving program and see what distances your range finder is calculating.

The lookup table provided by the distance calculator program will give you values that are fairly close to reality but to get the highest precision, you should carefully aim your range finder at known distances and generate your own lookup table manually using the actual measurements. This will correct for misalignment and lens distortion.

Conclusion

Having a laser range finder on your robot can allow you to dart from place to place at full speed without worrying too much about running into anything. If you were to mount it onto a hobby servo to swivel it around, you could quickly create a map of where your robot can




Figure 12. Self portrait taken with the linear image sensor.

and cannot go. Your robot could navigate from room to room fairly easily if it simply scanned each room for gaps in walls that were the correct width for a door. You could do this with a high degree of confidence since the distances reported by a laser range finder have a fairly high degree of accuracy using sub-pixel calculations, even at longer distances. What could you do with a laser range finder on your robot? **SV**

RESOURCES

Jameco Electronics
www.jameco.com

Mouser Electronics
www.mouser.com
 Sells the TSL3301 chip.

Edmund Optics
www.edmundoptics.com/US/
 Sells lenses.

Custom Computer Services, Inc.
www.ccsinfo.com
 Sells the C compiler used for the PIC code on SERVO's website.

Borland
www.borland.com/us
 Sells the C++ compiler used for the PC code on SERVO's website.

Spark Fun Electronics
www.sparkfun.com
 Sells a laser module that you could use for your range finder.

Back to school? Build a robot!

Serial 8-Servo Controllers

Pololu's serial servo controllers are compact and high-performance, featuring 0.5-microsecond pulse resolution and individual speed and range control for each channel. Both versions are available fully assembled or as partial kits that require only the connectors to be soldered in.



Serial 8-Servo Controller
 Features DB9 connector for direct connection to a PC serial port.
 #0728 partial kit: **\$23.95**
 #0727 fully assembled: **\$26.95**



Micro Serial Servo Controller
 With a 0.91" x 0.91" outline, this controller fits just about anywhere!
 #0208 partial kit: **\$17.95**
 #0207 fully assembled: **\$19.95**

Motor Controller Specials!

Using low-voltage, high-current toy motors is easy with the compact Low-Voltage Dual Serial Motor Controller. With a motor supply voltage range of 0-7 V, you can run low-voltage motors off of one cell, and you won't have to give up any power thanks to low on-resistance MOSFETs that deliver up to 5 A per channel. **Low-Voltage Dual Serial Motor Controller #0120: \$31.95**

Low-Voltage DSMC and Double Gearbox combo #0670: \$39.95



Tamiya 70168 Double Gearbox
 Two independent gearboxes in one compact package! #0114: **\$8.95**

See our web site for more specials!

Call 1-877-7-POLOLU or visit www.pololu.com for other robotics and electronics solutions.

Pololu
 Robotics & Electronics

6000 S. Eastern Ave. 12D, Las Vegas, NV 89119