The following paper was published in the March 2000 issue of EOM (Earth Observation Magazine), 4901 East Dry Creek Road, Suite 170, Littleton, CO., USA, 80122.

THE LOW DOWN ON LIDAR

By Robert A. Fowler

Every once in a while a technology suddenly seems to catch peoples' imagination and you start hearing about it all the time. That's been happening with LIDAR recently. What is LIDAR? Well, like radar, it's an acronym, except in this case it stands for LIght Detection And Ranging. What does that mean? It is the technology, which uses light, specifically a laser light, to measure distance.

Now you might say that's pretty old hat and there's nothing new about it, and you would be almost right. People have been using the electromagnetic spectrum to measure distances for nigh over fifty years. But over the last few years recent developments in associated fields have made this method of obtaining distances from aircraft platforms much more reliable.

It's an interesting conundrum of physics that nothing travels faster than light. It is (as far as we know) the phenomena that travels at the ultimate speed in our universe. Contrary to popular science fiction, even if you could approach the speed of light, you would be unable to actually catch up with it. Indeed, the laws of special relativity have essentially proven (strangely enough) that even if you did approach light speed, the relative speed of light to your speed remains the same, i.e: light still travels to or from you at the same speed, 186,000 miles per second, no matter how fast you are travelling.

Regardless, special relativity at aircraft speeds on this planet is a non issue, as it essentially has no effect on our distance measuring capability of light (unless you get down to Planck distances,* and we don't have any way of measuring them anyway)!

So what has been happening that suddenly makes airborne LIDAR hit the news? As I said, it's the associated technologies. Laser systems, even airborne laser systems have been around for a long time. The accuracy that they can measure to has been in the centimeter range for almost as long. But absolute accuracies in that range have been true only of lasers fixed on the ground. Once you place a laser system onto a moving platform – the position of the platform becomes the limiting factor.

Until Global Positioning System satellites were launched there was no way of knowing exactly where a moving object was in relation to any ground based coordinate system. Dead reckoning, using speed and direction detection devices were also fraught with limiting factors. Even fairly sophisticated inertial technology was only marginally successful in determining absolute position.

The Inertial Part

Let's look at what inertial systems do. Inertial systems (often referred to as IMUs – inertial measurement units) use another law of physics which says that objects spinning at a very high rate tend to keep their relative position in space. In other words if you take something like a child's toy top and manage to make it spin incredibly fast, it will not topple over.

Inertial systems essentially then consist of a minimum of a couple of gyroscopes, each of which spins a globular mass within a gimbal or cage. To reduce the friction of such devices, today these spinning objects are electro-magnetically held in suspension.   One of the gyros will be oriented in a vertical direction and a second fixed at right angles to it. Because of the wonderful effects of gravity, spinning objects in the near vertical will orient themselves to true vertical. This phenomenon provides us with a relative datum to work with.   Thus, as we know the axis of our spinning object is straight up and down, if the object on which the device is located tilts over (as happens when you bank during an aircraft turn) the gimbal cage turns with the aircraft while the spinning mass remains vertical. This allows the measurement of the angle of tilt.

If we couple this effect with another sophisticated device known as an accelerometer, we can start to measure velocity effects.   To imagine an accelerometer, think of a cardboard box with a string suspending a weight hanging from somewhere near the center top of the box.   Then, when you give the box a shove, the box moves forward ahead of the weight. The difference in distance the box moves before the weight moves and catches up with it is a function of its acceleration. If we can measure the difference in shift and have a very high accuracy timing device, we can calculate the measurement of acceleration and then distance.   Coupled with the gyroscope, our inertial system can now tell us how far, how fast and in what direction we are moving, relative to the point where we started.

While the basic theory of this is actually very simple, getting it all to work with accuracy is not. Inertial systems are hamstrung by two major limitations. One is how accurately you can make the measurements of tilt and acceleration (physically how accurately the angle changes can be measured and the arcs described), and two, how accurately you can measure time. Naturally, the smaller you make the spinning object and its cage, and the smaller you make your accelerometer; the more critical these small measurements become.   So, as with any technical equipment, miniaturization brings its own manufacturing problems.

Now let us assume, and rightly, that manufacturers have developed sophisticated inertial systems which work extremely well and measure minute effects incredibly accurately. However, once we place this equipment in an aircraft, it is subject to a host of other effects. For example, we know we can bank an aircraft so extremely that the effect of gravity is temporarily suspended and anything not tied down is going to float. However, even at lesser banks, there are very small effects that upset the inertial system. The result is, with these combined effects of motion, acceleration, turns and so on, after a couple of hours errors will accumulate and the inertial system begins to lose its sense of quite where it is. After several hours the system will accumulate errors that start to multiply exponentially.

At some point before this happens we need to be able to tell the system "whoa, we're not there - we are here." This is where developments in GPS have come in.

The GPS Part

Again, Global Positioning Systems have been around for something over a dozen years now, during which time manufacturers have gradually made amazing improvements to overcome the degrading of the signals by the military. As a result, in the last four years or so, with specific equipment and software it has been possible to obtain the position of the GPS antenna on an aircraft, using translocation, to within a few centimeters. If the aircraft antenna is ideally situated directly above the place where the laser (or an aerial camera for that matter) is located and the position has been surveyed in accurately, then it becomes possible to create the conditions to provide highly accurate positioning and thus aerial surveying. (Incidentally, the reason for the location of the GPS antenna above the sensing equipment is to reduce the effects of crab. Pilots can generally correct for the other rotations but crab is the most difficult rotation for the pilot to overcome and can reduce GPS accuracy significantly if you are stupid enough to put the GPS antenna on the tail or some other remote place.)

I am going to assume everyone knows roughly that GPS works by using essentially a simple resection of radio signals from a number of satellites whose orbits and positions at any one instant in time are accounted for. As well, using these same satellites to record the same data at a receiver on a fixed, known station on the ground will allow a final tie in a three dimensional triangulation to the coordinate system of choice. (Well it's simple in theory – again, doing it isn't quite so easy!)

OK, you are saying, the GPS can correct the IMU for position, but what about the rotations? One of the nice things about gyroscopes is even if we slide around the corners a bit too fast, once the pilot straightens up on line, the spinning axis of the gyroscope will gradually settle down to its vertical position, again through the pull of gravity. What we can't account for, though, is the "slippage' in the accelerometers, unless we have a GPS to every so often correct the position.

The Laser Part

Laser, is yet another acronym for Light Amplification by Stimulated Emission of Radiation. That isn't quite as horrible as it sounds and simply means that by passing a specific frequency of light through a prism you concentrate the beam and frequency so that a relatively weak light source can travel a long way with relatively little distension.

LIDAR systems generate a laser pulse which then, depending on what it hits, is reflected back (or partly absorbed) and the return spurt of light is recorded on a sensor. The light emitting diode transmits these pulses of light extremely rapidly, but the speed of light is such that the receiver sensor senses the return “echo” before the next pulse is sent. The time it takes for the signal to travel down and back multiplied by the speed of light and divided by two (because it goes there and back) is the distance to the ground (or object) from the laser head.

However, at this stage that is simply a distance, and without the other two components it means little. But, if we know the position of the laserhead in coordinated space (from the GPS) and its attitude (which way it is tilting from the IMU) then we can calculate a position and elevation using relatively simple rules of geometry, and be reasonably sure we are right.

The result of all of these developments coming together is the increased confidence in airborne LIDAR accuracy. Depending on the type of system in use, we can say under good conditions that the laser will generate results accurate to a couple of centimeters. The inertial systems can generate attitude to 0.1degree, which changes the positional accuracy depending on the flying height (because while the angular error doesn’t change, the positional error gets worse the higher you fly). Airborne GPS properly installed, used and processed will yield results in the 5-10 centimeter range. When you add all of these components together, while on occasion they might cancel each other out, the reality is you will get a guaranteed accuracy of around 15 centimeters or six inches. Anyone who tells you they can do better (with current technology, anyway) doesn’t understand what they are talking about.

Not all LIDARs are the same

There are a number of different LIDAR systems made by different manufacturers. Each system has its unique characteristics. As a rule of thumb, LIDARs made to work over water do not work as well over land and vice versa. The reason is, to penetrate water a specific wavelength of laser is used (usually in the green end of the spectrum). To obtain good reflections over land, LIDAR are produced with wavelengths in the infrared/near infrared frequencies. If a LIDAR is made to cover large areas at high level flying heights, unless the characteristics are changed it will be too powerful to use on low level flights. It might just burn out the sensor or be a hazard to things below. (Eye safe use is a significant concern with some LIDAR systems.) And vice versa. A system manufactured to operate at low flying heights will not have the power to work at greater heights above ground. Each type of LIDAR has its special applications and use.

On the other hand, in spite of anything you may hear to the contrary, a system that flies higher will not be quite as accurate as a system that flies lower. Now that may be arguable, as the laser part remains remarkably accurate over a wide range of distance. However, there will be some generalization of signal partly as a result of the footprint of the beam. But, the major degradation in accuracy is a function of the IMU. Any miniscule angular measurement errors are simply magnified and the computed easting and northing position of where the beam hits the ground has a larger circle of error the higher you fly. That is simple geometry that you just can't change.

Recorded LIDAR data, of course, are relatively unintelligent. The machine doesn't know whether it is recording reflections from the ground, a building, a leaf of a tree or some other feature.   Research is underway using systems which will record intensity of return signal. This is a feature that laser manufacturers are starting to push. At the moment, the practicality of intensity information is vague.   Intensity will provide the sort of information which on radar imagery is called "texture."

From this texture it is hoped eventually that software will be developed which will ascertain what the reflection is from: the type of vegetation, rock, sand, asphalt or a variety of other surfaces, depending on the strength of signal returned. The reality is, no one with a commercial unit, to my knowledge, is really using the intensity values on production projects and there is no laser-specific software to make the use of these data intuitive. So intensity values may remain a university project for the next little while or until someone develops software that makes them really useful.   Regardless, post-processing data is and will remain a part of the equation and we will briefly touch on that later.

Hopefully, by now you can see airborne LIDARs are technically extremely complex bits of equipment. In fact, they are actually *three* very complex bits of equipment that everyone hopes will work together.   There is the GPS which works with a very high precision clock, an IMU which uses another very high precision clock and a laser which needs another high precision clock. Anyone who has operated a LIDAR will tell you they are not easy to look after and are subject to various quirks in general, and odd system-specific quirks for each individual unit! Learning how to work one and cope with the quirks usually takes a pretty long time.

Nevertheless, in expert hands they can be amazingly accurate and amazingly efficient methods of obtaining ground and feature information.

The Processing Part

Each manufacturer supplies their equipment with software of varying capability and functionality. This allows, on a fundamental level, the removal of irrelevant data.

Obviously, if a LIDAR is flying over jungle it will receive returns from leaves at the top of the tree canopy. Some signals will return from leaves and branches between the canopy and ground, and some will inevitably (and somewhat surprisingly) be returned from the ground. Of course, depending on how thick the tree canopy is (and I have been in jungle where a lot of the time you simply cannot see the sky), there will always be places where, because of the rapid firing, the laser beam will manage to miss everything in between and actually hit the ground and get a return. This is one of the unique advantages of the system that no other mapping technology can duplicate. For, while in theory aerial photography can also “see” to the ground if the camera is directly above a small hole in the canopy, the adjacent photo, which allows the photogrammetrist to see in stereo and calculate the elevation, is taken from a different position. As a result the photogrammetrist can only see in stereo at the top of the canopy.

In a nutshell, in tree covered areas the manufacturer’s software will take the first returns and assume with something like 99.99% certainty that this represents the top of the tree canopy and then the lowest returns and assume that these are the real ground. Regular features such as buildings, which could fall between the two, usually have such a distinctive footprint that they are immediately obvious to the person processing the data. To aid in such interpretation, almost all LIDAR systems come equipped with a visual tracking system, mostly a video record, though it could be frame camera. These are “event located” using the GPS signal, so that the operator isn’t scrolling through miles of videotape trying to figure out where he is.

There are a few companies (very few) who have tied a digital frame camera to their LIDAR, and can provide orthorectified imagery as a direct result. This, it should be mentioned, creates a whole new series of problems. [It is not a matter of bolting the frame camera next to the laser and using simple geometry to make corrections to the pixels. There are huge technical and computational obstacles to overcome – but these are perhaps the subjects of another article.] Nevertheless, this creates a valuable additional deliverable for certain applications. And that simple phrase also sums up a lot of what LIDAR can do.

For some will tell you LIDAR will do anything and everything for you. I prefer to ignore the hype and ask people to look upon LIDAR as another tool or sensor which will help solve your, or a client’s specific problems. Like satellite imagery, radar surveys, photogrammetry or ground surveys, LIDAR has the applications it suits well and where it can be economically worth while. It can provide very rapid, very accurate surveys for specific situations, and we will examine a few of these in an irregular series over the next few months.

Robert Fowler, O.L.S., C.S.T., C.E.T. writes irregularly for EOM and for a number of other publications. He is manager of sales and marketing for Lasermap Image Plus and can be contacted at: bobf@lasermap.com

**Planck distance 10^{-33} centimeter*