Panoramic 3D-Imaging Using Single-Photon Counting Laser Radar

Markus Henriksson*, Julia Hedborg, Per Jonsson, Lars Sjöqvist
Swedish Defence Research Agency (FOI)
Box 1165, 581 11 Linköping
SWEDEN

*mahe@foi.se

ABSTRACT

The high data rate of single-photon counting array detectors for laser radar (LADAR) applications now makes it possible to use 3D systems for surveillance. A mast-mounted system could acquire 3D-data with centimeter accuracy and decimeter level resolution at up to ca 1 km range of the full 360 degree environment around a camp or a temporary encampment with update rates of several times per minute. Model experiments of single-photon counting LADAR measurement towards forest edges have been performed. The experiments show excellent resolution and several tens of meters penetration into the forest. Scaling of scan rates and signal levels show that a sensor array can be scanned at several revolutions per minute while maintaining sufficient signal levels for 3D-imaging. The vertical coverage will be set by the array size, where the currently available detector size allows 128 lines using a field of view of one to two degrees. Because of the stationary mounting of the sensor, all target surfaces are measured from one direction only, and every angular pixel can be processed separately as a simple histogram over range allowing parallelization and short lag times. Change detection then allows recognition of suspicious activity for further study. This application of novel LADAR technology thus could help to protect military installations against unwanted observers and planning of assaults.

1.0 INTRODUCTION

1.1 Scenario

At all types of military positions, permanent bases and facilities as well as more temporary encampments, it is important to discover any suspicious activities in the surrounding areas, including e.g. unwanted observers or preparations for an attack. Surveillance cameras, in the visible or IR, can observe these areas and by algorithms automatically detect and flag movement. 2D-cameras, however, have a very hard time detecting persons or objects hidden inside the edge of a forest. Another problem with camera surveillance is that a very large number of pixels is needed to cover the complete surroundings with high enough resolution to discover a person at several hundred meters range. If cameras are mounted in continuously rotating sensor turrets, to cover the scene with high resolution at regular or irregular intervals, the 2D information may not be sufficient to detect a change from the previous passage of that exact position.

To solve this problem we suggest the use of a high resolution 3D LADAR sensor. The sensor would be mounted on a rotating turret and sweep the surroundings to create a panoramic 3D view. A laser radar based 3D-imaging system can potentially penetrate up to hundreds of meters into forests, and certainly reach any position from which it would be possible to observe the sensor position itself. A single panoramic 3D image would, however, contain too much data for efficient real-time detection of suspicious activity. However, if the sensor is rotating around a fixed axis, every time it is looking in the same direction it should see the same objects at the same distance. It would thus by change detection in 3D be possible to discover any moved objects between two passages. This scheme shows promising capabilities to autonomously, and with low false alarm rate, detect activity that needs to be inspected by an operator.

1.2 Photon counting laser radar

Photon counting laser radar, at least in the sense discussed in this paper, uses a detector that triggers on the first photon that enters the detector after it has been armed. The most popular and relevant detector type is the Geiger mode Avalanche Photo Diode (GmAPD), also designated Single-Photon Avalanche Detector (SPAD). This detector type is biased above breakdown voltage for the semiconductor material, so that a single absorbed photon will bring the detector to saturation by inducing an avalanche. After detection the bias voltage is reduced in order to stop the avalanche, a process called quenching, and then raised again to re-arm the detector for detection of another photon. For the distances discussed in this paper arming occurs once per laser pulse, possibly delayed compared to the pulse so that a minimum detection distance is created. For one laser pulse the system will give a very accurate measurement of the time until the first detected photon, or the information that no photons were detected during the measurement gate. For a single measurement there is no possibility to separate laser photons reflected from the target from background light or detector dark counts. When multiple measurements are collected in a histogram (distributed in time bins) the target returns will pile up at the bin corresponding to the distance of the target, while noise sources are randomly distributed in time. The histogram thus corresponds to the full waveform distance measurement with Poisson-distributed noise.

One major advantage of photon counting laser radar is that the temporal (distance) resolution can be very high as the detector contribution is in principle set by the accuracy of measuring the flank of the detector pulse, and not by the analog bandwidth. In our application this allows the separation of camouflage, in the form of vegetation, and targets. Using a 20 ps laser pulse and a single pixel InGaAs SPAD overlapping surfaces separated by 5 cm can be detected separately [1]. This is approximately a factor 10 better than for a system using linear mode APDs. SPAD arrays, on the other hand, are currently limited by the time bin size, and need 20 cm separation in range between surfaces to properly detect both.

Another advantage is that the high sensitivity and histogram collection where the detection probability can be far below 1 per pulse allows the use of low pulse energy lasers at high repetition rate. High repetition rate lasers with low pulse energy generally have lower SWaP (size, weight and power) than high pulse energy low repetition rate lasers used in linear mode flash 3D laser radar. The disadvantage of using many laser pulses and photon counting detectors is obviously that an integration time is needed to image a target, which is a disadvantage when considering rapidly moving targets.

The factor in photon counting arrays that make the surveillance application possible is the high data rate. Compared to linear mode flash array LADAR detectors that run at 10 to 100 Hz frame rate a photon counting array with frame rate on the order of 100 kHz can cover a much larger field of regard despite the need for multiple measurements in each direction. Depending on the need for dynamic range, multiple return performance, range accuracy and reflectivity measurement the number of measurements required in each direction can vary considerably. Except for the case when accurate measurements of reflectance are needed the efficiency of photon counting technology is expected to be superior regarding area coverage and total laser pulse energy per pixel [2].

There are reports of linear mode photon counting detectors that have single-photon sensitivity and still are able to detect multiple photons per pixel per laser pulse [3]. If these become available with similar specifications for data rate as the SPAD arrays and with good properties in other ways, the performance of the sensing concept discussed here could be further improved.

2.0 EXPERIMENTAL PANORAMIC 3D IMAGING USING SINGLE PIXEL SPAD

To illustrate the feasibility of change detection using a photon counting sensor, measurements of a forested scene with and without targets were collected with a scanning single pixel photon counting laser radar system. Target positions were both open and hidden behind vegetation.

2.1 Measurement system and scene

The measurement system was constructed at FOI using a picosecond erbium fiber laser and a single pixel InGaAs SPAD. A monostatic transceiver with Cassegrain optics is placed on a pan-tilt table to allow scanning across the scene. The system has approximately 160 μrad lateral resolution and is continually swept in horizontal direction to cover the scene row by row. The FWHM distribution in range direction for measurement of a point target is 392 ps, corresponding to 59 mm, which can be taken as a measure of the range resolution. The data collection is limited to 4096 bins, which together with the selected bin size of 64 ps limited the scene depth to 39.3 m. More details of the measurement system are available in [1].

Targets were placed at the edge of a forest, approximately 260 m from the sensor position, and 101 sweeps over 5° laterally were performed at different vertical angles. The measurement was performed three times, twice with the targets present and once without any manmade targets. As shown in Figure 1 some of the targets are very visible, while others are more hidden. The sweep velocity was set so that a single measurement of the full scene containing 101 sweeps required approximately 10 minutes.

Figure 1. Photograph of the scene with targets that is used for illustration of the sensor concept.

2.2 Signal processing and experimental results

The raw data is a large number of detections with angle-angle-range information. The vertical angles have 101 discrete values, but the horizontal angles are continuously varying. The range data is given as one of 4096 time bins. As the data was collected by a SPAD every detection is only a range value with no intensity information. Based on the number of detections compared to the number of transmitted laser pulses it is improbable that a detection corresponds to more than a single photon except for when hitting a corner cube reflector at one position in the scene. Thus we also have no significant problems with reduced average detection rates at the back of the measurement gate because of detector saturation.

The first step taken in the analysis is to group the detections in discrete ranges of horizontal angles to create a 2D image where every pixel contains a histogram of the frequency for different range bin values in that direction. The horizontal pixel size was here chosen as approximately half of the lateral resolution to limit the loss of resolution from the angle binning. On average every histogram contained around 160 detections at different ranges. For the pixels covering the corner cube there are approximately 4000 detections in the histogram, giving an estimate of 4 % detection probability for the scene. The detection probability was limited by laser power. A thorough effort to separate background and noise from signal has not been performed, but an estimate is that on average 25 of the 160 detections in each histogram are a result of detector noise and background light, while the rest are reflected laser photons.

To reduce the influence of shot noise in the histograms a low pass matched filter is applied. The matched filter is the time-reversal of the instrument response function as measured against a point target. The filtered histograms for one pixel are shown in Figure 2. At 273.9 m there is a peak in two of the histograms belonging to one of the manmade targets in the scene. In the third dataset collected without the targets present this peak is as expected missing, and instead there is a peak at 274.8 m from an object (a tree or the ground) further back. This introduction of a peak and removal of one further back are the signs of something changing in the scene that we are looking for. To complicate the problem we see that peaks between 258 and 263 m that are reflections from trees in front of the target are not constant, and will not fully cancel when subtracting the histograms. The difference is caused both by shot noise and by leaves and branches moving in the wind to alter the ideal laser radar response.

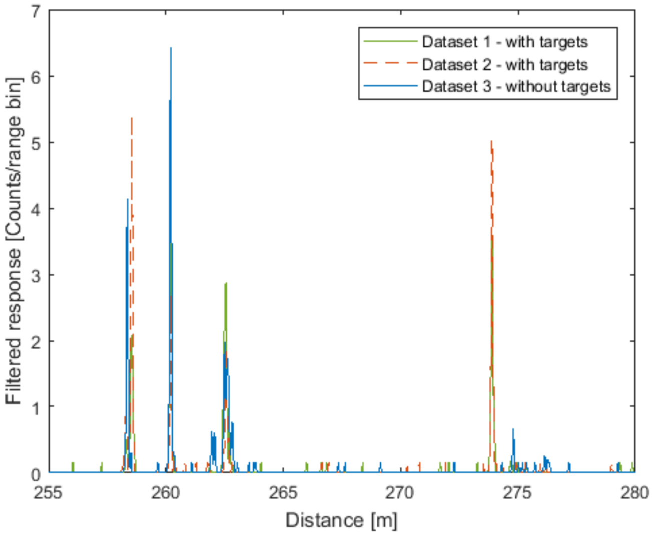

Figure 2. Histogram of three measurements in for the same transverse pixel in the panoramic image. The responses from 258 to 263 m are from vegetation, while the peak at 273.9 m is from a manmade target that was present in two measurements but not the third. In the dataset without target a reflection from the ground is present at 274.8 m, which was blocked by the target in the other two datasets.

The inability of single histograms to discriminate between new targets and natural movement of vegetation forces the use of full 3D data for discrimination of changes that are significant. It is difficult to simultaneously study the histograms representing the full 3D volume, and thus we present the data for one single row in Figure 3. The graphs show the right part of the scene from Figure 1, and have also been cropped in distance to make the objects in the image large enough to be visible. By comparing the images it is easy to see the three targets with straight lines, and also the missing signal at longer ranges behind these targets. This visual inspection of 2D views would, however, be very impractical in operative use.

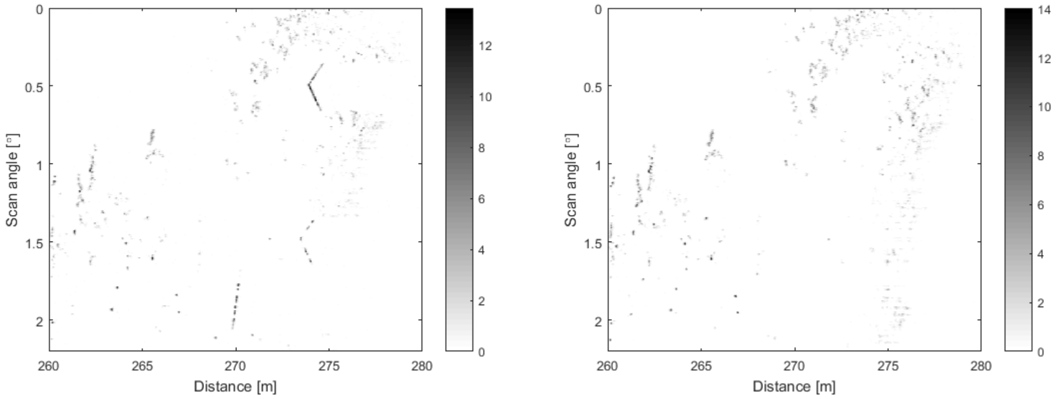

Figure 3. Low-pass filtered response for part of one slice at constant vertical angle through dataset. The left figure is with the targets present and the right without the targets. The colour shows intensity (photons returned per unit of time).

Because the sensor is rotating around a fixed axis the histograms from different datasets correspond to exactly the same lines and can be directly subtracted. A view of the result is shown in Figure 4. The targets are visible as linear features in red, and the blocked background as areas of consistently blue-green groups of features behind them. Groups of mixed blue-green and yellow-red features are caused by the vegetation moving in the wind. This view increases the possibility to detect changes, but still needs an automatic detection. Our belief is that filtering can discriminate between areas with both positive and negative changes from small movements and more significant changes in the 3D-structure. Another attribute of data to use is that changes that are not the last detected target in a certain direction is unlikely to be a target of interest, as that corresponds to something that is smaller than the instantaneous field of view (IFOV) of each pixel.

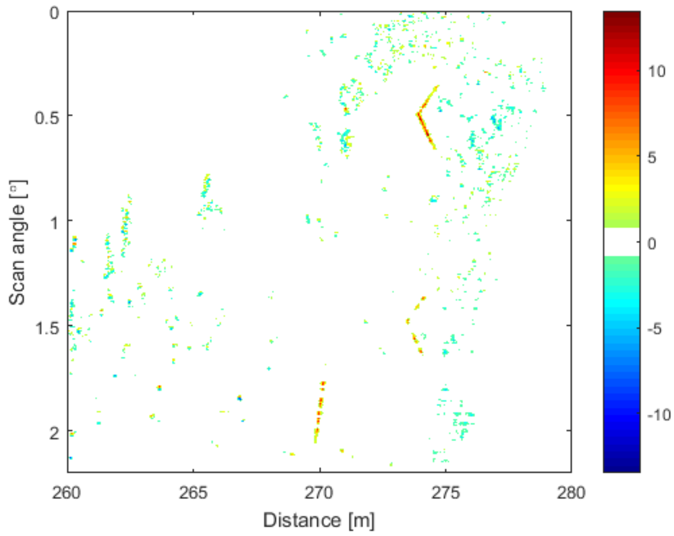

Figure 4. Difference of the two images in Figure 3. The targets are visible as linear features in yellow-red, and the blocked background as areas of consistently blue-green groups of features behind them. Groups of mixed colour features are caused by the vegetation moving in the wind.

The discussion so far has considered only single sweeps at fixed inclination angle. To view the whole scene at one instance we, however, need to move to a 3D view considering all the data at the same time. This further increases the need for automatic signal processing supporting the operator.

Because of the large scene the true returns were extracted from the histograms as all local maxima fulfilling certain thresholds in amplitude and width. The data points were converted to point clouds for visualization in a commercial 3D lidar software. The point clouds calculated from the measured data without and with the targets present are shown in Figure 5 and Figure 6, respectively. The large amount of details captured by the high resolution 3D sensor makes finding the differences difficult, even if a few should be obvious.

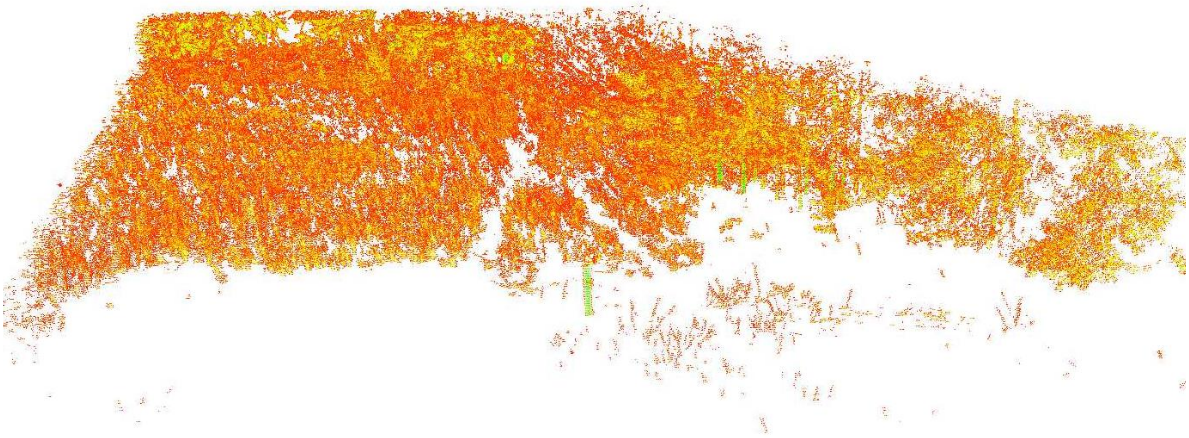

Figure 5. Point cloud of scene without the targets present. Colour goes from red for weak signals, via yellow to green for strong signals. The view is from higher angle than the sensor position.

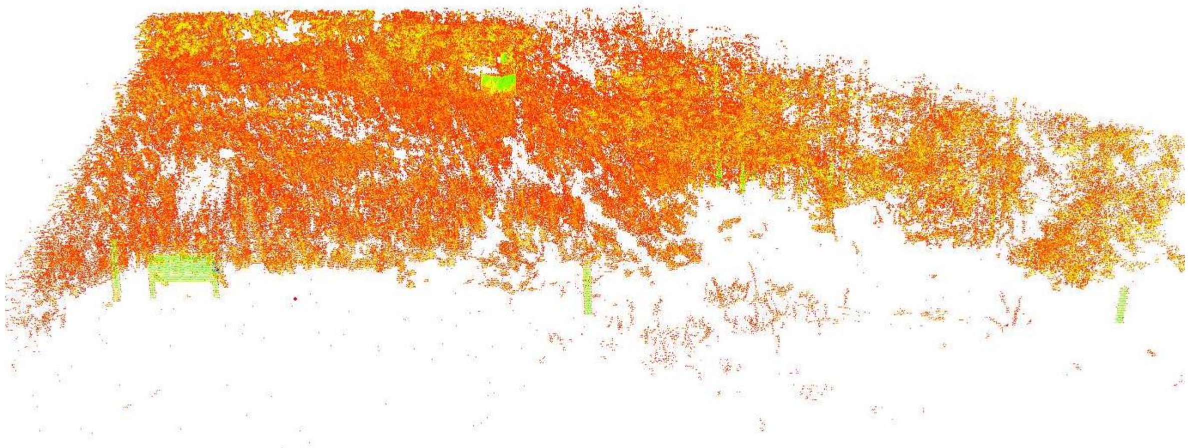

Figure 6. Point cloud of scene with the targets present. Colour goes from red for weak signals, via yellow to green for strong signals. The view is from higher angle than the sensor position and the same as in Figure 5.

There are several available methods to detect changes between the measurements [4]. Here we only perform a simple test to show that it is possible. For the point cloud under test we first filter out only the last point detected from every histogram, as we assume that targets are opaque. Finally we compare the remaining points to the reference point cloud. If the intensity of the point is much larger (here set as factor 2) than the sum of all points in the reference point cloud in a 0.2 m cube surrounding this point it is counted as a new point. The view of this is shown in Figure 7.

Additionally we perform the opposite test where we look for silhouettes of objects that have disappeared from the reference point cloud. Again we reduce the data to the last point in each direction to save computation time. Finally we compare the point intensity value to the total intensity in a small cube at the same position in the point cloud with targets in the same way as before. In this way objects that are no longer visible because a new target has blocked the view of them can be visualized and show a silhouette of the target. The obscured objects shown in Figure 8 are good indicators to in which directions targets have been added, and are especially useful for low reflectance targets. Because the background here is quite rough with high semi-dry grass and other vegetation the silhouettes are not clear enough to show the shape of the added objects, but the increased density of points indicate that something has happened.

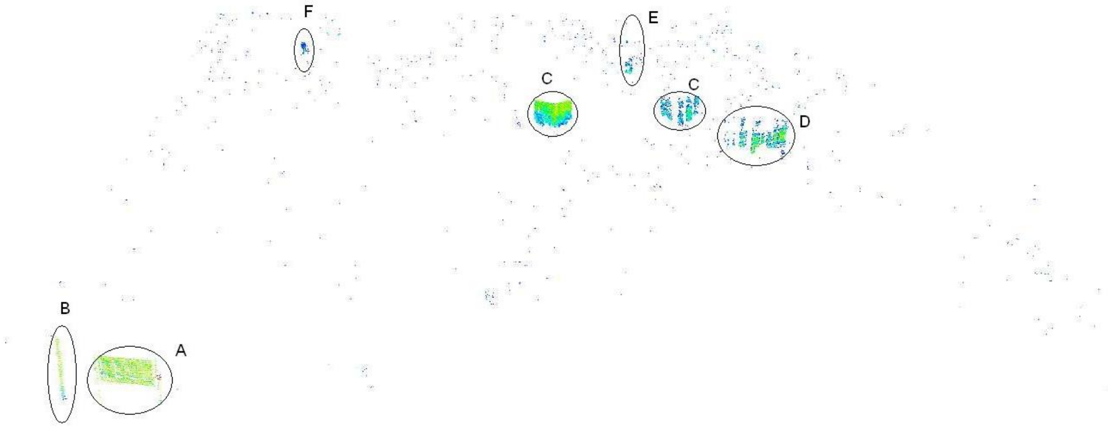

Figure 7. Points that are present in scene with targets, but not in scene without targets, when a simple processing is used. Ellipses show detected changes and annotations are explained in the text.

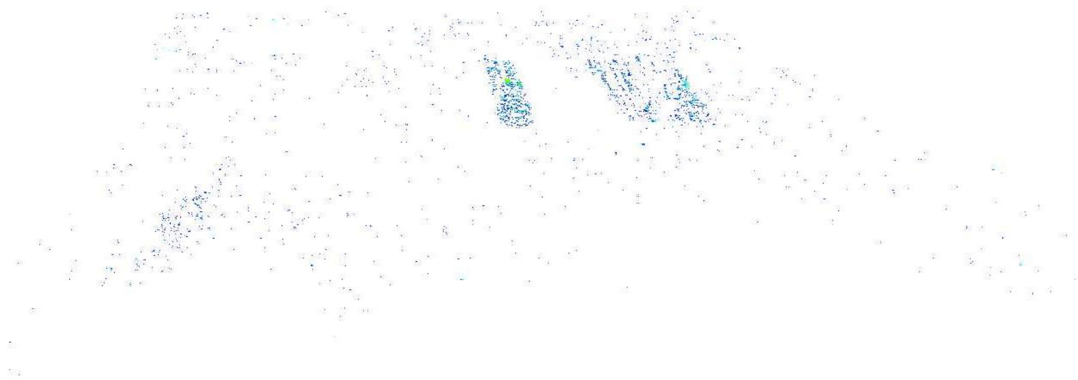

Figure 8. Points that were present in scene without targets, but have disappeared when the targets were added. The view is the same as in Figure 7.

The change detection in Figure 7 shows all intentional changes between the two measurements, but also a scattering of false alarms that are not removed with the simple methods used for the analysis in this paper. More advanced signal processing methods can likely yield even better results. The detected changed targets were:

A. One structured target for 3D-sensor characterization in the front left of view

B. Two vertical poles at the front of the gate, one is outside the shown view in Figure 7

C. Two large corner structures in the middle of the gate, one in open view and one behind some trees

D. A larger flat surface with patches of varying reflectance behind slightly denser vegetation blocking the view of especially the upper part (that also had the darker colors)

E. A mannequin quite far into the forest (sideways between the two corners). For the mannequin only a small part, mainly of the legs, is visible. For the torso and the head the system did not manage to penetrate the foliage of the relatively dense young aspen trees in front of it with the used sweep rate and signal processing.

In the back left of Figure 7 there is an additional group of dense points (F) that does not correspond to an intentional change. This group is just at the back of the measurement gate, and may correspond to a branch that has moved across this threshold by changing winds. All other false alarms are scattered single points that should be easy to remove with further signal processing.

3.0 EXPERIMENTAL SPAD ARRAY MEASUREMENTS

A measurement system using a SPAD array, the Princeton Lightwave Inc. (PLI) Falcon [5] is under development at FOI, but measurement data for panoramic imaging suitable for change detection is not yet available. The system uses a 100 mm focal length optics, together with the 50 µm pixel pitch giving 0.5 mrad spatial pixel resolution. The 128×32 pixels are illuminated by a Keopsys fiber laser with 0.57 ns pulse length with the beam profile expanded to cover the rectangular field of view by cylindrical lenses. The camera and laser is mounted on a rotation stage to allow horizontal scanning for panoramic imaging, with synchronization of angle by fork sensors. The angular synchronization of the rotation stage and the camera needs to be improved in the future as the current solution relies on the assumption of constant angular rotation rate. Another necessary future improvement is data transfer at unlimited rotation using high capacity slip rings. Currently cables are limiting the rotation to less than one revolution.

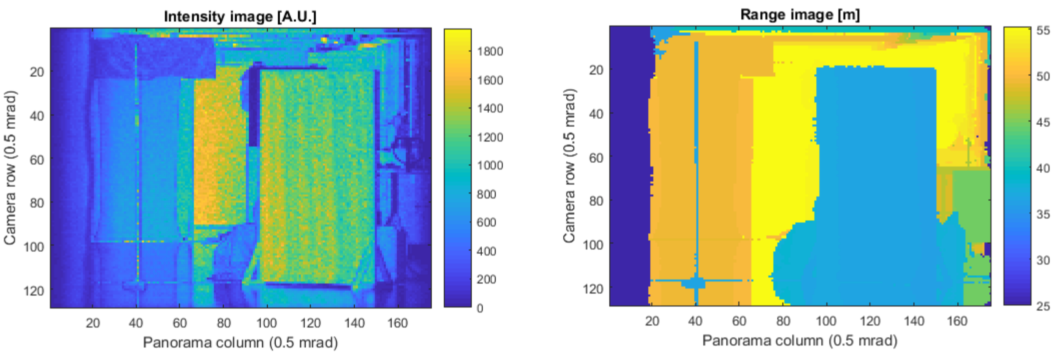

Figure 9. Example 4° FOV panorama from a 3D imaging LADAR using a photon counting array detector under development at FOI. The data was acquired at an angular velocity of 20°/s with 0.5 mrad spatial resolution.

An example panoramic 3D image recorded at an angular rotation rate of 20°/s, limited by the rotation stage, is shown in Figure 9. The scene as acquired indoors with a maximum range of ca 50 m. This is currently ongoing work and improvements are needed in many ways for the system functionality, data processing and result presentation.

4.0 DISCUSSION

Current generation commercially available SPAD cameras, here represented by PLI Falcon, have 128×32 pixels and a maximum measurement frame rate of 90 kHz. This is by no means a limit, but we will use this value for our calculation. This corresponds to 369 million range measurements per second, but to allow the measurement of multiple return waveforms the probability of detecting any photon in each measurement has to be kept at around 0.5 or lower.

We assume that we want a lateral resolution of 0.2 m at 400 m distance, corresponding to 0.5 mrad IFOV, to be able to classify a target as human sized and thus interesting to investigate further. If we assume that the sensor is placed at 10 m above ground and we should cover the span from the flat ground at 100 m out to horizontal view at 10 m above ground this corresponds to 100 mrad (5.7°), and could with quadratic IFOV be covered by 200 pixels. With the PLI Falcon placed with the 128 pixels vertically we would thus need two sensors with their FOV stacked in height. Design of a single chip sensor with the necessary number of pixels in the vertical direction is expected to be possible if there is sufficient demand. If UAVs should be detected significantly more pixels would be needed as a much larger vertical angular span would need to be covered. For UAV detection other operating modes not optimized for high range resolution penetration of vegetation are probably also more suitable, and this problem will not be covered in this paper.

The full rotation is 6.28 radians, and would correspond to 12566 IFOVs. For every laser pulse the system performs 32 measurements, one per column, at each vertical position. A system rotating at 1 Hz, would then perform 229 measurements within the pixels of the panoramic image, with pixel sizes corresponding to the IFOV, per revolution. It should however be kept in mind that 229 measurements is not the same as 229 photons returned from targets at different ranges. Because of the need to avoid saturation from the closest object in every direction a significant portion of the measurements will detect no photons. There will also be some number of measurements that are triggered by background light or detector dark counts. For optimum adjustment approximately 50 % of the measurements may contain a target return, but because of varying reflectance and distance in the scene the proportion will vary between pixels. 229 measurement would thus likely be sufficient to measure the distance to the strongest reflecting surface in every direction, but not to discover targets hiding in a forest edge. By reducing the update speed to once every 10 s the number of measurements per pixel is increased to 2292 per pixel, which would allow more detailed multiple return processing where every histogram contains around 1000 detections. A 0.1 Hz update rate (6 rpm) is similar to many radar systems, and deemed to be sufficient

The measurement range R may vary between 100 and ca 1000 m, even though we optimized the lateral resolution for 400 m. Neglecting absorption and scattering the laser radar response from an extended target (larger than the IFOV) varies as R^{-2}. A difference in distance of a factor 10 would then produce a factor 100 difference in signal level. To avoid saturation at the shortest distance the laser power would thus be so low that the number of measurements that detect a target return at the longest distance is very low. But because the measurement is performed repeatedly it would be possible to adapt the laser power to the measurement distance in every direction. Based on statistics from previous revolutions the laser power could be varied by rotation angle. It is also probable that the distance varies significantly with the vertical angle. One approach to vary the laser power along the array would be to use a master oscillator power amplifier array (MOPAA). In a MOPAA the light from a single laser is split in several fibers, each containing a separate amplifier. Every fiber would have a separate collimator and the beams stacked vertically to together fill the rectangular field of view of the sensor. In this way the laser power in different sectors of the sensor FOV can be varied independently to fit the target distance and reflectivity in that direction. By using a wide bandwidth laser and not matching the lengths of the different amplifier chains the beam combination will be kept incoherent and avoid any interference effects. The use of an amplifier array would also reduce any issues with component damage or laser safety as the different channels are not overlapping at the exit aperture.

There are situations where the envisioned system would not work. One is close to any strong light source, as the rising or setting sun, where the sensor would be saturated by the high background level. Normal daylight illumination can be handled with high F-number optics or neutral density filters, but may then have to be compensated with higher laser power to maintain optimum photon detection probability. Another problem is that strong retro-reflections may damage the current generation of sensors. As this is not a problem for single pixel SPADs it should be possible to solve this issue in the future. A retroreflector would however still hide anything at a longer distance in the same direction by saturating the detector, and the surroundings should be cleaned from retro-reflecting objects for optimum performance. The size of the object that is to be detected may be a problem. If only the head of an observer is protruding from behind a solid obscuration this may with the suggested resolution only be visible in a single pixel in the panorama. The skin would give very low signal for a system with 1.55 µm laser wavelength, but there should be a significant negative signal of disappearing background. The problem would be to reliably separate this from natural background variations caused by shot noise and vegetation moving in the wind. A single sensor could also not reliably protect a large facility as it will not see through solid obscurations, from behind which it may still be possible to observe or target a different part of the compound.

5.0 CONCLUSION

We have shown, through experimental evidence and simple calculations, that photon counting 3D laser radar on a rotating platform can be used as a surveillance sensor to detect adversaries moving into position for observation of the sensor position. The advantages compared to 2D cameras are primarily in forested environments where the time-of-flight 3D laser radar can penetrate denser vegetation, although some small part of the IFOV still needs to provide free line-of-sight for the target to be visible. The use of photon counting laser radar, as compared to classical scanning or flash 3D LADAR systems, is motivated by better range resolution and higher data rates. Objects that are large compared to the lateral resolution should, based on the image examples, be easy to automatically detect, while smaller objects while giving a significant signal may be hard to separate from natural variation in the background signal caused by e.g. vegetation shifting in the wind. Calculations show that resolution for detecting human-sized objects at up to approximately 400 m and an update rate of 0.1 Hz for the full 360° perimeter should be possible with the sensor performance of currently commercially available photon counting sensors.

ACKNOWLEDGEMENTS

We acknowledge support from Mr. Lars Allard with measurement system construction and data collection and from Mr. Folke Berglund, Mr. Roland Lindell and Mr. Bengt Ericson with setting up the scene. This work was performed in projects supported by the Swedish Armed Forces.

REFERENCES

[1] M. Henriksson, H. Larsson, C. Grönwall, and G. Tolt, "Continuously scanning time-correlated single-photon-counting single-pixel 3-D lidar," *Opt. Eng.*, **56(**3), 031204 (2017).
[2] P. F. McManamon, P. Banks, J. Beck, D. G. Fried, A. S. Huntington, and E. A. Watson, "Comparison of flash lidar detector options," *Opt. Eng.*, **56(**3), 031223 (2017).
[3] J. D. Beck, R. Scritchfield, P. Mitra, W. W. Sullivan, A. D. Gleckler, R. Strittmatter, and R. J. Martin, "Linear mode photon counting with the noiseless gain HgCdTe e-avalanche photodiode," *Opt. Eng.*, **53(**8), 81905 (2014).

[4] R. Qin, J. Tian, and P. Reinartz, “3D change detection – Approaches and applications,” ISPRS J. Photogramm. Remote Sens. **122**, 41–56 (2016).

[5] http://www.princetonlightwave.com/wp-content/uploads/2017/01/PLI-Falcon.pdf