Three-dimensional imaging laser radar with a photon-counting avalanche photodiode array and microchip laser

Marius A. Albota, Richard M. Heinrichs, David G. Kocher, Daniel G. Fouche, Brian E. Player, Michael E. O'Brien, Brian F. Aull, John J. Zayhowski, James Mooney, Berton C. Willard, and Robert R. Carlson

We have developed a three-dimensional imaging laser radar featuring 3-cm range resolution and single-photon sensitivity. This prototype direct-detection laser radar employs compact, all-solid-state technology for the laser and detector array. The source is a Nd:YAG microchip laser that is diode pumped, passively *Q*-switched, and frequency doubled. The detector is a gated, passively quenched, two-dimensional array of silicon avalanche photodiodes operating in Geiger mode. After describing the system in detail, we present a three-dimensional image, derive performance characteristics, and discuss our plans for future imaging three-dimensional laser radars.

OCIS codes: 280.3420, 280.3640, 120.0280, 110.6880.

1. Introduction

The direct-detection three-dimensional (3-D) laser radar[1] (ladar) concept is straightforward. Light from a short-pulse laser illuminates a scene of interest. The light reflected from this scene is imaged onto a two-dimensional (2-D) grid of detectors. Rather than measuring intensity, as in a charge-coupled device (CCD) or a conventional 2-D camera, these detectors measure the photon time of arrival, and thus the range. With each pixel coded with range, the ladar produces an angle–angle-range, or 3-D, image.

We have developed a 3-D ladar with single-photon sensitivity, 3-cm-range (depth) precision, adjustable angular resolution, and 128 × 128 pixels. The illuminator is a Lincoln Laboratory-built, diode-laser-pumped, passively *Q*-switched microchip laser operating at 532 nm.[2] The electro-optical receiver is a Lincoln Laboratory-built 4 × 4 array of avalanche photodiodes (APDs) that was passively quenched and gated to operate in Geiger, or photon-counting, mode. Each of the 16 APD pixels in the detector array is mated to custom external timing circuitry. The field of view (FOV) of the entire 4 × 4 APD array is scanned to generate larger images with up to 128 × 128 pixels. The entire system is robust and can fit into a small van to be taken to field sites for experiments. This ladar is a prototype for a more advanced 3-D imager that Lincoln Laboratory is developing that will use larger arrays of Geiger-mode APDs with integrated complementary metal-oxide semiconductor timing circuitry.[3]

In Section 2 we describe the laser source and APD array with associated detector board electronics, the timing circuitry, the optical system components, and the process of data collection and rendering. In Section 3 we show an example 3-D image collected with this prototype ladar. In Section 4 we use data from the prototype to determine its performance characteristics. In Section 5 we discuss our future efforts to develop advanced and compact 3-D ladars. We conclude by summarizing the results and examining the potential applications of 3-D imaging laser radars.

The authors are with the Lincoln Laboratory, Massachusetts Institute of Technology, 244 Wood Street, Lexington, Massachusetts 02420-9108. When this research was performed, M. A. Albota, R. M. Heinrichs, D. G. Kocher, D. G. Fouche, B. E. Player, M. O'Brien, J. Mooney, B. C. Willard, and R. R. Carlson were with the Laser and Sensor Applications Group. M. A. Albota (albota@ll.mit.edu) is now with the Optical Communications Technology Group. B. F. Aull is with the Advanced Imaging Technology Group. J. J. Zayhowski is with the Quantum Electronics Group.

Received 29 March 2002; revised manuscript received 8 October 2002.

0003-6935/02/367671-08$15.00/0

2. Enabling Technologies and Three-Dimensional Ladar System Components

A. Microchip Lasers

The Lincoln Laboratory-built, passively Q-switched microchip lasers are constructed by a short piece of Nd:YAG (gain medium) diffusion bonded to a similar piece of Cr^{4+}:YAG (saturable absorber).[2,4] The pump-side face is coated to transmit the 808-nm pump light and reflect the laser light at 1064 nm. The output pulses have a transform-limited frequency spectrum, fundamental transverse mode with diffraction-limited divergence, and linear polarization. The high intensity of the output beam allows its frequency to be doubled when a small piece of KTP is placed near its output face. The energy efficiency of frequency doubling to 532 nm is typically 50–60%.[4–6]

For use in this 3-D ladar prototype, we selected a doubled-Nd:YAG microchip (4 cm long, 1.5 cm in diameter, hence the name) laser that generates 250 ps (FWHM) pulses at 532 nm. This laser is pumped with 1.2 W at 808 nm from a commercially available diode laser (SDL Inc., Model 2372). The microchip laser is operated at a repetition rate of 1 kHz by external modulation of the diode pump-laser current, and the output energy per pulse is 3 μJ. This relatively low energy-per-pulse illumination source is suitable to image short-range (< 1-km) targets. In general, the laser power requirements for a ladar imager will depend on several parameters, including target range and reflectivity, the two-way propagation path loss, the receive aperture area, and on the detection efficiency of the receiver. Microchip lasers with output energy per pulse over an order of magnitude larger than the one used in these experiments have been developed at Lincoln Laboratory. Although the data presented herein were collected with a 3-μJ/pulse laser, we have more recently collected 3-D images at ranges beyond 1 km with a 30-μJ/pulse source.

B. Geiger-Mode Avalanche Photodiode Arrays

The APDs are front-illuminated and uncoated Lincoln Laboratory-built devices microfabricated in a p-i-p-i-n geometry in silicon.[3] The array consists of 16 circular detectors arranged in a 4 × 4 grid. The 30-μm-diameter active-area APD elements are separated by a 100-μm pitch for a 7% fill factor. The APDs are reverse biased and gated above the breakdown voltage to operate in photon-counting mode, in which an electron–hole pair generated by the absorption of a single photon initiates an avalanche process causing the APD to break down. In response to a single photoelectron, Geiger-mode APDs yield a fast, high-amplitude electrical pulse that can be used to trigger high-precision computer-automated measurement and control (CAMAC) timing circuitry. We measured breakdown voltages of the order of −41 V at a temperature of 10 °C, normalized dark-count rates of 3–10 kHz, and timing jitters of less than 290 ps with 250-ps-wide laser pulses on 30-μm-diameter APDs. The APD array used in this 3-D ladar is a front-illuminated 4 × 4 device without antireflection coating and has a measured quantum efficiency of ~20% at 532 nm.

C. Avalanche Photodiode Detector Board Electronics

The 4 × 4 APD array is mounted on a five-layer preamplifier circuit board. The preamplifier electronics diagram for each element in the array is shown in Fig. 1. The board is populated on both sides, with two ground planes sandwiching microstrip transmission lines for the gate and wide-bandwidth (WB) APD signals. Two outputs are available on the preamplifier board for each detector in the array. The narrow-bandwidth (NB) output is used to monitor the activity on individual detectors and also to adjust the individual detector bias voltages with the bias trim, should there be any nonuniformity in the breakdown voltage across the array. The WB output is used to measure precisely the photon arrival time. Referring to the circuit diagram of Fig. 1, a common negative dc bias is applied to all elements of the array on a common anode connection. The dc bias is −39 V, which is 2 V lower in magnitude than the required breakdown voltage. When Geiger-mode operation is desired, the cathode voltage is gated to a positive voltage large enough to result in the desired voltage drop across the device. In photon-counting operation, the APD is typically gated on with +8 V (6-V overbias) for 2 μs every time the laser fires (every 1 ms). The gate voltage is applied through 100-kΩ resistors to each detector cathode to limit the average current to safe values at all times. To achieve rapid turn-on times, a 47-pF speed-up capacitor was added across the 100-kΩ resistor. The voltage on the cathode is buffered for off-board use by a NB, or low-frequency, high-input-impedance operational amplifier stage. This operational amplifier has a unity-gain bandwidth of 4 MHz, a 9-V/μs slew rate, while also having extremely low-quiescent-power dissipation to minimize the amount of heat dissipated on the board. The voltage pulse created on the cathode by the avalanche signal is coupled by a 2-pF coupling capacitor into a 50-Ω coaxial cable, which connects to off-board amplification and the timing discriminators. Such a small capacitance driving a 50-Ω load effectively impresses the time derivative of the APD cathode avalanche waveform on the coaxial cable, resulting in a signal that consists primarily of the high-frequency (timing) content of the waveform. The resulting avalanche pulses are approximately 10 ns wide and have fast rise times of less than 2 ns.

D. Avalanche Photodiode Timing Circuitry

The WB avalanche pulses from each pixel on the detector board are sent to CAMAC modules through coaxial cables. These timing signals are cascade amplified by two stages of 20-dB voltage gain (Phillips Scientific Model 776), recorded by CAMAC-based LeCroy electronics and sent to a computer through a general-purpose interface bus module for processing.

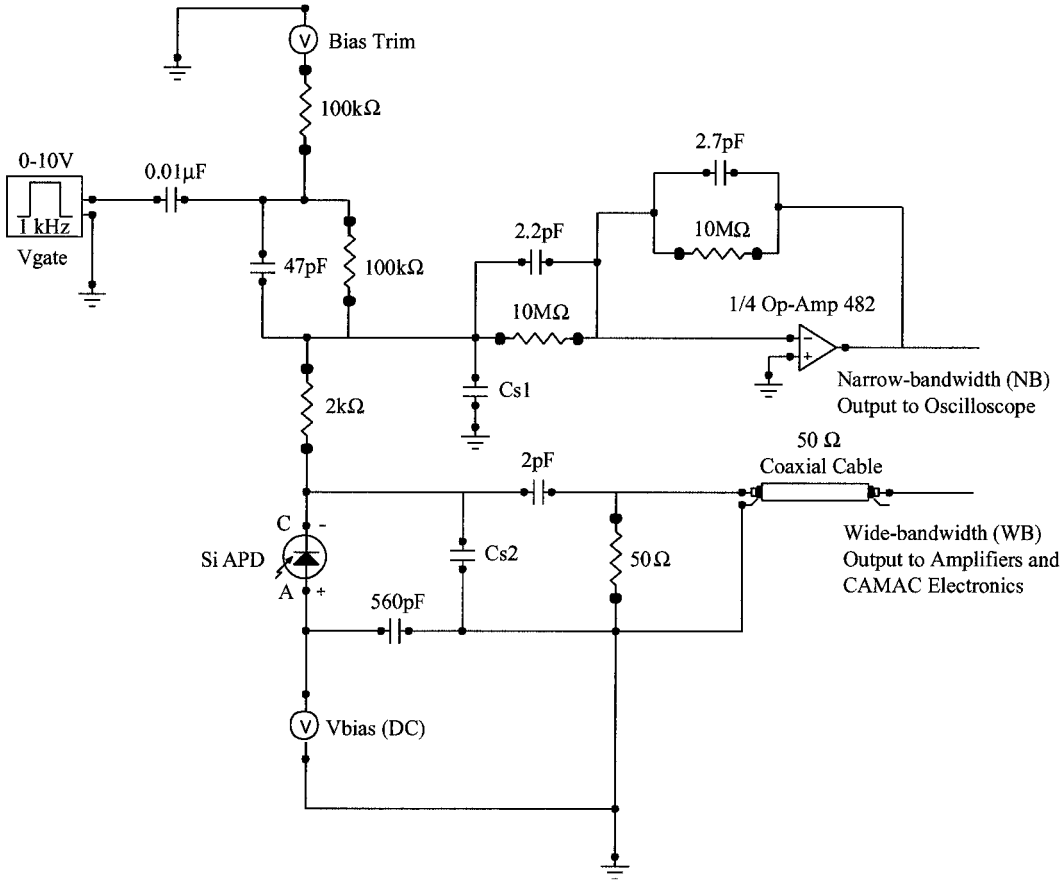

Fig. 1. Preamplifier circuit diagram for each of the 16 detectors in the 4 × 4 APD array. Cs1 and Cs2, measured stray capacitances (Cs1 + Cs2 = 3.25 pF); op-amp, operational amplifier (see text).

The time-to-digital converters digitize time with a least-significant bit of 75 ps. A measurement cycle begins with the custom circuitry issuing position commands to the scanning mirrors and modulating the pump laser to full power. Microseconds later, the laser pulse is emitted and sensed by a wideband detector (EG&G FND100) to record the transmit time. The transmit pulse is turned into a logic signal by a LeCroy 3412E discriminator and sent to a LeCroy 4303 and 4300 transmit time-to-digital converter. The receive window for data collection, typically 117 ns for all 16 channels, is started by the range clock. The receive time-to-digital converter is stopped by the individual wideband APD output. The range of a target is derived when the transmit time, the receive time, and the offset of the range clock are combined. Each measurement is transferred to LeCroy 4302 digital memory modules with the emitter-coupled logic bus. The two memory modules are ping-ponged so that data from one half of the frame are captured in each module. The full memory is unloaded by an IEEE 488 interface to the PC such that data recording is continuous.

E. Three-Dimensional Ladar Optics

Figure 2 shows the functional diagram of the 3-D ladar optical head. The 3-μJ linearly polarized laser pulses pass through variable diverger optics and then through a λ/2 wave plate for polarization adjustments. The fundamental YAG wavelength at 1064 nm and the residual pump beam at 808 nm are

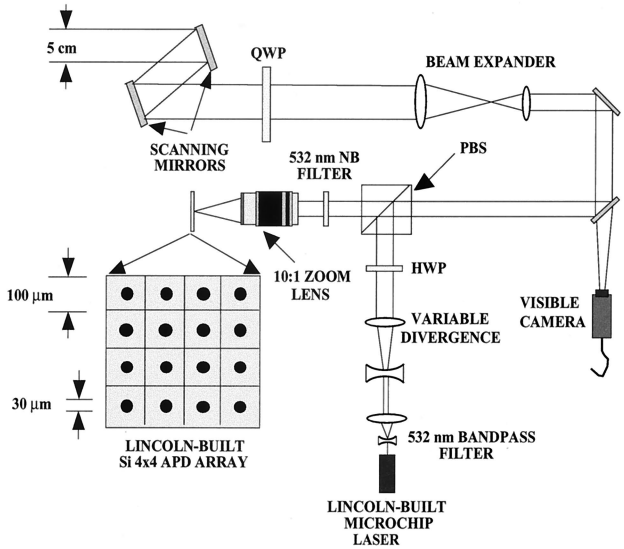

Fig. 2. Optical layout of the 3-D laser radar. QWP, quarter-wave plate; PBS, polarizing beam splitter; HWP, half-wave plate.

blocked by a bandpass filter (Andover 775FW82-25) with 95% transmission at 532 nm. The beam then passes through a polarizing beam splitter and a $\lambda/4$ plate acting as a transmit–receive switch. The system operates in a monostatic configuration, using the same 5-cm-diameter aperture for both transmit and receive beams. The instantaneous field of regard (IFOR) of the 4×4 detector array and the divergence of the transmitted laser beam are separately adjustable. Laser divergence is set to match the IFOR of the 4×4 array of detectors for efficient collection of laser photons. The total FOV of the system is built up when we scan the IFOR of the 4×4 detector in both azimuth and elevation. A pair of single-axis scanning mirrors (Cambridge Technologies, Model 6900) raster scans the IFOR of the 4×4 array over the angular extent of the scene of interest and provides frames of data with up to 128×128 pixels. The scan rates and laser-pulse rate are synchronized so that the image of the 4×4 array moves over one IFOR between pulses. Therefore a 128×128 pixel frame is composed of 32×32 patches of 4×4 pixels, each patch illuminated with a single laser pulse. The light backscattered from the target reflects from the same scan mirrors and is directed toward the detector by the polarizing beam splitter. The photons returning from the target are imaged by a 10:1 zoom lens onto the focal-plane array after passing through a narrow-bandpass filter (1-nm FWHM) (Andover 532FS02-25) with ~58% peak transmission at a 532.2-nm center wavelength (~55% transmission at 532 nm). The frame rates for 32×32 and 128×128 pixel images are 4.5 and 0.6 Hz, respectively, limited only by the speed of the mirror scan motors. We select angular resolution or sampling density by adjusting the divergence of the laser illuminator and receiver IFOR.

F. Three-Dimensional Ladar Operation and Data Collection

The operational situation for this 3-D imaging device first involves the transmission of a temporally short laser pulse by the laser radar to illuminate a target. At some time, before photons reflected by the target are expected to arrive back at the detector, the overbias gate voltage is turned on and the individual timers begin counting. We keep the dark-count rates at low and constant levels by cooling the APD array to 10 °C. The receive window (range gate), which is the recording interval beginning when the APDs are overbiased to Geiger mode, is typically 117 ns long and starts at a selectable time after the laser pulse is detected in the transmit beam optical path. This range gate can be scanned in depth to find a target at an unknown range or to collect background (free-space) photons from regions in front of or behind a target. As the backscattered photons arrive, they are imaged onto the APD detector array. These photons generate avalanches in the photodiodes, which trigger the timing electronics to stop counting. The count values reached by each of the timing elements correspond to the relative range of the corresponding part of the target. This array of ranges represents the 3-D surface profile of the target. Background noise is minimized through a combination of narrowband spectral filtering (the 1% bandwidth of the filter is 3.6 nm) and temporal filtering by short-duration gated-mode operation.

Having a range for each pixel, we can display the 3-D image in a variety of ways. The first is to assign a different color to each range and apply the appropriate color to each pixel in the displayed 2-D array. A second is to render a 3-D model from the data and display the model as if it were photographed. We use the commercially available software package IDL (Interactive Data Language) for color-coded pixel and 3-D rendering. A third is to create a point cloud and project these points onto a plane representing the viewing screen. In a point cloud, each pixel is assigned a point with x and y coordinates corresponding to the pixel position in the array and a z coordinate corresponding to its range. We use the software package Igor Pro to display and rotate point clouds. The human eye–brain system can better recognize the 3-D nature of the point-cloud image when the projection plane is rotated progressively, as if the observer is moving around the recorded image, or equivalently as if the image is rotating in front of the observer.

3. Example Image from the Three-Dimensional Ladar Prototype

We collected 3-D images of various vehicles, planes, boats, wires, people, foliage, buildings, and other objects with this ladar prototype. In Fig. 3 we show as an example an image of a 1987 Chevrolet Astro van. This target was selected because it offered a wide reflectivity range, from the absorbing black tires, asphalt, and side mirrors to the highly reflecting license plate. The 128×128 pixel data were collected in broad daylight from a *single* vantage point situated approximately 60 m in front of the vehicle. The transverse (x and y axes) and depth (z axis) resolutions for this image are approximately 2–3 cm. The detectors were biased into Geiger mode approximately 50 ns before the first photons reflected from the van arrived. Range is encoded in the gray scale, but not linearly, as shown by the accompanying bar. White indicates that there was no detector firing for that pixel. The four versions of the image differ in the number of data frames processed. One of the 16 detectors was nonresponsive; we set its value to that of one of its nearest neighbors.

The single-frame image (upper left) shows the van well, even though it has a few extraneous points caused by firings from dark current and background light and has several pixels that did not fire. To reduce noise points and fill in blank pixels, we processed multiple frames and looked for range coincidences. That is, for a given pixel, we looked at the set of recorded ranges from multiple pulses and chose as the range to display that value about which the most points were clustered. This process works well because firings from background light and dark cur-

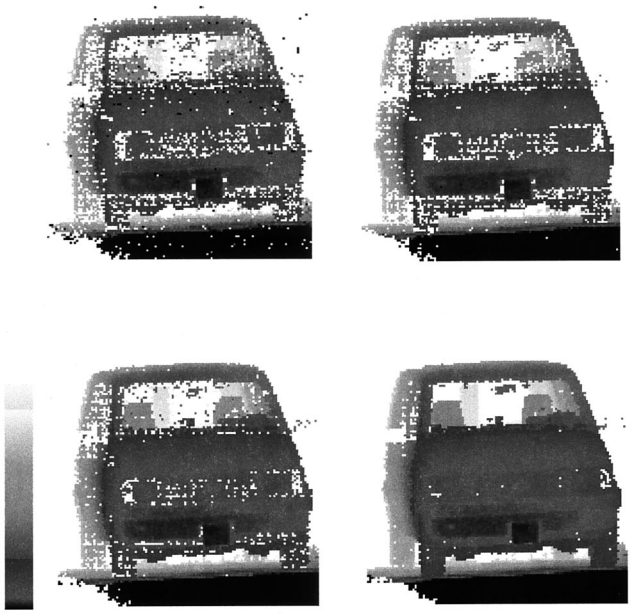

Fig. 3. Images of a Chevrolet Astro van obtained with the 3-D ladar prototype. The single-frame image (upper left) clearly shows the features of the van. The other three versions result from range-coincidence processing on 3 (upper-right), 10 (lower-left), and 110 (lower-right) frames. Range is encoded in the gray scale, with the nearest ranges being darker. The 3-D images were collected from a *single* vantage point, 60 m in front of the vehicle, in midday light on a sunny day. The range resolution is approximately 3 cm (see text).

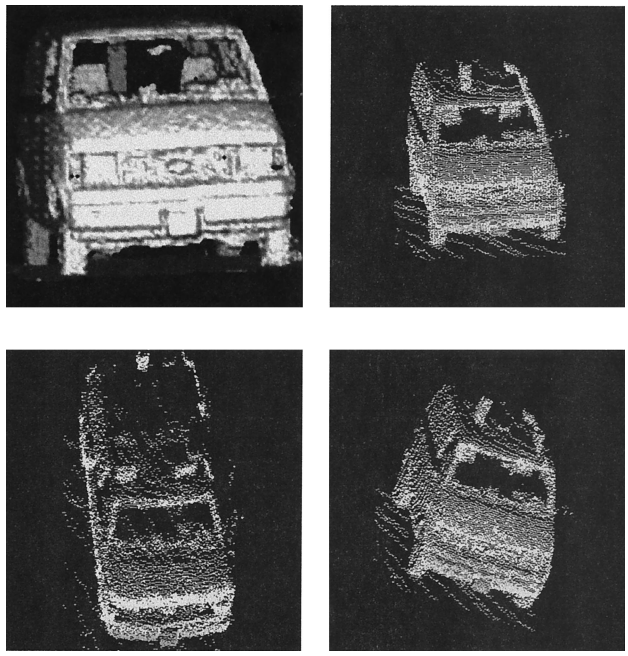

Fig. 4. Image of the Chevy van computed from range-coincidence processing on 110 frames and shown in other formats. In the upper left is a 3-D model rendered from the angle–angle-range data. The other three renditions are point clouds viewed after various amounts of computer rotation. Rotation of the image in software better reveals shapes, sizes, and relative positions of different parts of the van.

rent occur at ranges randomly dispersed throughout the gate, whereas firings from the laser reflection occur only at or near the range of the van element imaged to that pixel. The other three versions of the van image shown in Fig. 3 result from range-coincidence processing on 3 (upper-right), 10 (lower-left), and 110 (lower-right) frames. The number of blank pixels and noise points decreases as the number of frames increases.

Figure 4 shows the 110-frame version in other formats. At the upper left is a surface-rendered version created from the data points. The other three versions are point clouds after computer rotation by different amounts. This single image of the van is much more graphic when viewed dynamically as it is rotated by a computer.

4. Performance Characteristics of the Three-Dimensional Ladar Prototype

A. Poisson Statistics

In Geiger mode, the first primary electron that is created initiates the current surge that triggers the timing circuitry. The creation of primary electrons is a Poisson process. Therefore, if the photon rate at the detector r is constant, the time of the firing T is expected to be a random variable with an exponential distribution: $P(T = t) = r \exp(-rt)$. We verified this expectation using data from our ladar prototype. The data were previously recorded for another purpose. For these data, the laser was not used and the ladar receiver passively imaged an indoor scene that was flood illuminated by two 1-kW incandescent lamps providing a nearly constant photon rate. The narrow-bandpass filter in front of the detectors was removed to allow more light to fall on the detectors. The detectors were biased into Geiger mode (armed) 1000 times/s, and each time the data were recorded for the first 117 ns after arming. We recorded 125 frames of the scene, each with 128×128 pixels.

Our first step in data processing was to eliminate the complicating effects of detector-to-detector cross talk (discussed below). We did this by retaining, for each arming, only the earliest firing time of the 15 good detectors. This is equivalent to our treating the 15 detectors as a single detector (macrodetector) because the sum of Poisson-distributed random variables is a Poisson-distributed random variable. The electron-creation rate of the macrodetector is equal to the sum of those of the 16 individual detectors. Our next step was to select a subset of the 32×32 patches in the scene with the criterion that they all have nearly the same electron-creation rate. Then we could treat the results from individual patches in the subset as being drawn from a single Poisson process with a given rate. We wanted to process data from more than one patch to improve statistical averaging. We ended up choosing 243 patches for which the electron-creation rate varied by approximately $\pm 10\%$. This gave us 30,375 data points altogether

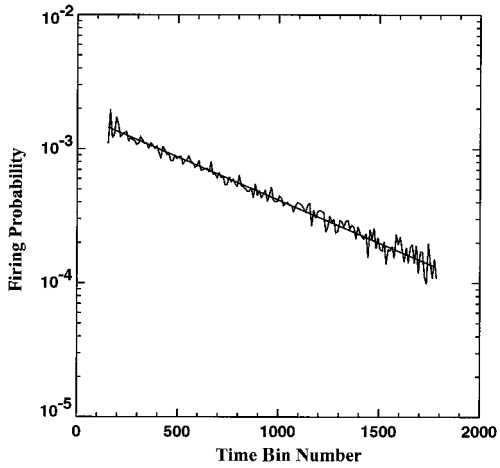

Fig. 5. Firing probability versus time bin number when the detector array is illuminated with incoherent light. The straight line is a least-squares fit to the data of this log-linear graph (see text). The exponential dependence indicates that firing of Geiger-mode APDs follows Poisson statistics.

(243 patches times 125 frames). Our final step was to calculate the fraction of the 30,375 points for which the macrodetector firing occurred in each time bin. We chose to divide the recorded times of firing into 0.75-ns intervals, or bins.

Figure 5 shows the results. The abscissa is the number of time bins after we armed the detectors. The ordinate is the fraction of data points occurring in a bin; this is equivalent to the probability of firing within a time bin. The straight line is a least-squares fit to the data of this log-linear graph. The fit is excellent, indicating that the firing times closely follow an exponential law. The slope and initial value of the straight line are 0.0147 and 0.0145, respectively, close to the equality expected for an exponential probability distribution. The amount of scatter of the data about the straight line is roughly consistent with what one would expect for a binomial distribution with the given number of trials (30,375) and probability of success (ordinate times 30,375). The extra scatter at the beginning of the curve is surmised to be caused by thermally generated electrons already in the absorption region of the detector when arming occurs.

Given Poisson statistics, the probability P that a detector fires in a time interval is 1-exp($-N$), where N is the average number of primary electrons created in that time interval. Thus N is equal to $-\ln(1 - P)$. To determine P, we bias the detector into Geiger mode m times ($m \gg 1$), count the number of firings f within the time interval of interest, and set P equal to f/m. P and N are nearly equal to one another when they are much smaller than one.

B. Dark Counts and Cross Talk

We measured the room-temperature dark-count rate of the Geiger-mode APDs of the prototype ladar. To do so, we used a black cloth to prevent light from reaching the detector array. Data from 200,000 armings of the 4 × 4 array were recorded. Each time, data were recorded for the first 117 ns after arming. We counted the fraction of the 200,000 armings that each detector fired within the 117-ns recording interval. The average value for the 15 good detectors is 0.0747%.

Some of these firings occurred not because of dark current, but because of cross talk occurring shortly after one of the other detectors fired. Cross talk occurs because photons emitted during the avalanche process can migrate to other detectors, where they can be absorbed and initiate an avalanche. To eliminate firings from cross talk from our calculations of dark counts, we count only the first firing of the detectors for each arming. This approach has a slight error, as it eliminates any firings from dark current after the first in each arming. The error is only 1% of the dark-count rate, because the probability of a firing of any kind from any detector was only 1% per arming, so that the probability of two or more firings from dark current on any arming (such firings are statistically independent of one another) could be no more than 0.0001. We find that the fraction of the time that a detector fired first, averaged over all detectors, is 0.0504%. This means that each dark-count firing created an average of (0.0747 − 0.0504)/0.0504 = 0.48 firings from cross talk within the 117-ns recording interval. The distribution of time delays between primary and cross talk firings has a peak between 10 and 20 ns and extends to the 117-ns limit set by the recording interval. It is reasonable to assume that the amount and time-delay characteristics of cross talk will be the same when the primary firings are initiated by photons rather than by dark current.

The 0.0504% average darkcount probability is for a 117-ns interval. Thus the darkcount rate is 4.31 kHz on average for the 15 detectors. The maximum and minimum dark count rates for the 15 detectors are 6.3 and 2.9 kHz, in good agreement with the values obtained from the APD characterization measurements reported in subsection 2.B.

C. Noise from Background Light

A Geiger-mode APD can be triggered not only by a laser photon reflected from the scene, but also by a sunlight or skylight photon reflected from the scene. We determined the background-light photoelectron rate for the Chevy van data as follows. For each pixel (except for those covering the street in front of the van), we counted the fraction of frames for which the detector fired during the initial part of the range window in front of the van. The fractions, or probabilities, for the various pixels were all much less than 1%. We reduced the probabilities by a factor of 1.48 to account for cross talk (Subsection 4.B). The probabilities are so low that they are equal to the average primary-electron rates (see end of Subsection 4.A). The pixel-to-pixel average value is 87,000 primary electrons per second. After subtracting the dark-count rate of 4310, we end up with a pixel-averaged value of 83,000 photoelectrons per second

from background light. This amounts to 0.0097 photoelectrons within our full, 117-ns range gate—a low amount of background-light noise.

The lighting conditions for the van data were noted only qualitatively. The Sun was at a zenith angle of roughly 40° and toward the driver's side of the van, so that most of the background light in the scene was from reflected skylight rather than reflected sunlight. There were no clouds. Suspended behind the van was a blue curtain of unknown reflectivity; it filled the FOV of the ladar. The reflectivity of the van was not measured.

We can roughly estimate the expected photoelectron-creation rate caused by reflected skylight and compare it with the measured value. We use the following values for the calculation. The spectral radiance of skylight reflected from a surface of 50% reflectivity is of the order of 2×10^{17} photons/(s m^2 srad nm) when the Sun is at 40° from zenith and at an azimuth angle near 90° and when the line of sight is nearly horizontal. For the van data, the bandpass filter was 1 nm wide, the instantaneous FOV for an individual detector was approximately 300 μrad, the receiving telescope had a diameter of 5 cm, the optics transmission for 532-nm unpolarized light was approximately 12%, the detector fill factor was 0.07, and the detector quantum efficiency at 532 nm was ~20%. Thus, for these conditions and for a resolved object, we expect reflected skylight to have amounted to an average of roughly 50,000 photoelectrons/s. This order-of-magnitude estimate is consistent with the pixel-averaged measured value of 83,000.

We also note that had the Sun been shining directly on the scene, reflected sunlight would have been roughly two to three times as strong as reflected skylight, or roughly 0.02–0.03 photoelectrons per 117-ns range gate. This is still a small value for background-light noise. The extra factor of 2–3 arises as follows. With the Sun at 40° from zenith, the spectral irradiance of sunlight incident on a horizontal or vertical surface is of the order of 1 W/(m^2 nm), or 3×10^{18} photons/(s m^2 nm) at a wavelength of 532 nm. Upon reflection at near zero angle of incidence from a diffuse surface of 50% reflectivity, the backreflected spectral radiance is of the order of 5×10^{17} photons/(s m^2 nm). This radiance is 2.5 times higher than that used above for reflected skylight.

D. Signal Level

To measure the signal level from the prototype, we illuminated a plate coated with barium sulfate, which constitutes a diffuse surface with near-unity reflectivity. The plate was at a range of 58 m and was highly resolved. The transmit power of the laser beam was attenuated by a filter with a nominal optical density of 2.4. The scan mirrors were turned off. The plate was illuminated approximately 100,000 times, and the fraction of time that each detector fired was computed. The average fraction is 0.215. The contributions to this from dark current and background light (room lights were off) are negligible. We reduce the average firing fraction by a factor of 1.48 to account for cross talk. Then we use the Poisson relationship to calculate the average number of photoelectrons created per laser pulse. It is 0.16 photoelectrons per pulse. The uncertainty, because of limited statistical sampling with these data and with the cross-talk data, is 0.01.

This value for the photoelectron rate is close to what would be expected from calculations by use of the radar equation for a resolved object. We use the following parameter values in the radar equation: The laser-pulse energy is 3 μJ, the attenuation by the 2.4 optical density filter is a factor of 251, the wavelength is 532 nm, the transmission of the transmitter optics is approximately 67%, the transmission of the receiver optics for unpolarized light is approximately 12%, the two-way atmospheric transmission is unity, the laser light is divided into 16 pixels with a fill factor of roughly 0.5, the diameter of the receiver telescope is 5 cm, the range is 58 m, the detector fill factor is 0.07, and the detector quantum efficiency is approximately 20% at 532 nm. For the calculation, we also assume a diffuse surface of unity reflectivity, so that the backscatter coefficient is $1/\pi/\text{srad}^{-1}$. Then the estimated signal level is 0.27 photoelectrons. This is within a factor of 2 of the measured value of 0.16. The agreement is good, considering that the estimate is uncertain by a factor of at least 2. The uncertainty is dominated by the attenuation of the optical filter. Not having measured the attenuation, we use the nominal value, and our experience with other highly attenuating filters of this type is that the nominal value can be off by as much as a factor of 2 in transmission. The other factor with significant uncertainty in our estimate is the fraction of laser light falling within the image of the detector array; we did not measure this and believe that our estimate could be off by 50%.

E. Range Precision

Range precision as used here is a measure of the pulse-to-pulse variation in the measured ranges to an object at fixed range. We also use the term range jitter. To determine range precision, we set up the prototype ladar to image a flat plate that is oriented in a plane perpendicular to the ladar line of sight. The scan mirrors were turned off, so that the image of the 4×4 detector array was stationary on the plate. We recorded the detector firing times for 23,000 laser pulses. We then calculated the FWHM of the distribution of recorded firing times for each individual detector. The detector-to-detector average FWHM is 200 ps, corresponding to 3.0 cm in range. The shapes of the time distributions are approximately Gaussian.

For these measurements, the detector-to-detector average probability of firing on a particular laser pulse was 99.3%. At such a high probability of firing, timing jitter is relatively small, for two reasons. First, a detector almost always fires in the early part of the reflected laser pulse, so that the width of the

laser pulse contributes little to the measured jitter. Second, variation in the times that a photoelectron first arrives at the detector's avalanche region decreases as the number of photoelectrons increases. For both reasons, range resolution at a low probability of detection is not as good as 3.0 cm.

5. Future Development

The prototype 3-D ladar demonstrated that 3-D images can be made with Geiger-mode APD arrays. Given this success with a 4×4 detector array and external timing circuitry, we have begun a program to develop the array technology further. Recently, we fabricated 32×32 Geiger-mode APD arrays and integrated them with commensurate arrays of digital timing circuitry. We are now testing a second-generation 3-D ladar using an integrated 32×32 focal-plane array. We expect that scaling up to larger arrays (128×128 pixels, say) will be relatively straightforward. This technology will allow for the capture of entire 3-D images on a single laser pulse at high frame rates.

To improve detection efficiency, we are exploring techniques for integrating a lenslet array with the detector array. Each lenslet focuses light at the focal-plane pixel onto the active area of the corresponding detector, effectively increasing the fill factor up to 100%. We can decrease the amount of cross talk by minimizing the capacitance of the preamplifier board and by building APD arrays with gaps between individual pixels.

Because the current Geiger-mode APDs are made of silicon and do not respond well to wavelengths beyond approximately 900 nm, we are developing APDs with III–V materials that will respond to near-IR wavelengths at 1064 nm and in the eye-safe regime around 1500 nm.

The energy per pulse and repetition rate (up to approximately 50 μJ at 10 kHz for 532 nm) from existing microchip lasers are sufficient for many applications. If an application requires more energy, it is straightforward to couple a microchip laser to a diode-pumped solid-state amplifier. For operation at eye-safe wavelengths ($\geq$1.5 μm), a microchip laser at 1064 nm can pump an optical parametric oscillator.[7]

With microchip lasers and compact microelectronic arrays of detectors and timing circuitry, one can envision building a 3-D ladar small enough to be man portable, yet capable of recording 3-D images with thousands of pixels at hundreds of frames per second. Data from multiple frames could be processed to see through obscurants, freeze motion in dynamic scenes, extend the range, or expand the FOV.

6. Conclusions and Potential Applications

We have developed a high-resolution, direct-detection 3-D imaging ladar using a microchip laser and silicon Geiger-mode APD array built by Lincoln Laboratory. This prototype imaging system was built to test and validate the interaction of the various ladar components and enabling technologies and also to gather high-resolution 3-D images. The microchip laser employed in this ladar utilizes diode pumping of Nd:YAG to generate 3 μJ of green energy per pulse when operated at a repetition rate of 1 kHz. The 4×4 APD array offers robust, solid-state, single-photon sensitivity. The ladar has generated 3-cm-resolution 3-D images of targets with a wide reflectivity range. These angle–angle-range images provide unambiguous geometric shape and orientation information, which can be of much greater value than conventional angle–angle–intensity data for many applications. The continuing development of these detector and laser technologies will lead to faster 3-D imaging systems with higher spatial resolution at longer ranges in compact packages.

The 3-D imaging ladar is well suited for a variety of applications, including the navigation of autonomous vehicles, robotic vision, remote industrial process monitoring, surveillance, and topography.

This research was sponsored by the U.S. Air Force under contract F19628-00-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the Department of Defense. M. Albota wishes to acknowledge useful discussions with Jonathan Bloch.

References

1. C. G. Bachman, *Laser Radar Systems and Techniques* (Artech House, Norwood, Mass., 1979).
2. J. J. Zayhowski and C. Dill III, "Diode-pumped passively *Q*-switched picosecond microchip lasers," Opt. Lett. **19,** 1427–1429 (1994).
3. B. F. Aull, A. H. Loomis, D. J. Young, R. M. Heinrichs, B. J. Felton, P. J. Daniels, and D. J. Landers, "Geiger-mode avalanche photodiodes for three-dimensional imaging," Lincoln Lab. J. **13,** 335–350 (2002).
4. J. J. Zayhowski, "Passively *Q*-switched microchip lasers and applications," Rev. Laser Eng. **29,** 841–846 (1998).
5. J. J. Zayhowski, "Ultraviolet generation with passively *Q*-switched microchip lasers," Opt. Lett. **21,** 588–590 (1996).
6. J. J. Zayhowski, "Ultraviolet generation with passively *Q*-switched microchip lasers: errata," Opt. Lett. **21,** 1618 (1996),
7. J. J. Zayhowski, "Periodically poled lithium niobate optical parametric amplifiers pumped by high-power passively *Q*-switched microchip lasers," Opt. Lett. **22,** 169–171 (1997).