Lab Notes

NEWS FROM AROUND LINCOLN LABORATORY

SATELLITE SAFETY

All Versus All Conjunctions

Weaving through a minefield of objects in Earth orbit is getting more difficult

You're driving along a proverbial springtime road in New England trying to avoid the potholes. You jerk the steering wheel right, then left, but eventually you still hit the deepest one because you can't quite figure exactly where your passenger-side tire is and the potholes come in very quick succession.

Now project yourself upward into the region of orbiting satellites travelling over 5 km/s. How do you avoid the oncoming "potholes" of other satellites and space debris? Arthur Lue and his associates in the Space Situational Awareness Group are concerned with the ever-increasing number of objects to avoid and how close they come to active satellites on a daily basis. They consider two issues—how many potholes there are and how accurately can they locate the passenger-side tire. If they can determine the precise location of every object in a satellite's orbit, it may not be necessary to jerk its steering wheel too often.

Lue worries about the expanding collection of space debris and satellites in orbit around the Earth. According to Lue, the current Spacetrack Catalog lists almost 15,000 objects greater than 10 cm in size. These include objects ranging from active and dead satellites and rocket bodies to misplaced tools from manned space flights to the myriad pieces of scrap metal resulting from the Chinese anti-satellite missile test of 2007 (about 2,500 objects) and the Iridium/Cosmos collision in 2009 (about 1,400 objects). "We're already near the Kessler Syndrome limit [the point at which there will be a runaway chain-reaction increase in the number of objects in orbit, [1]," Lue says.

Recalling the asteroid-Earth scenario of the movies *Armageddon* and *Space Cowboys*, Lue offers two alternatives. "Either you go out and destroy the asteroid or you predict its path accurately and move the people out of the city where it is going to hit." For satellites and debris, he offers the same two alternatives: get rid of all the debris or shift the position of the satellite to avoid the collision. A recent commercial on television shows the Air Force Space Command shifting the position of a

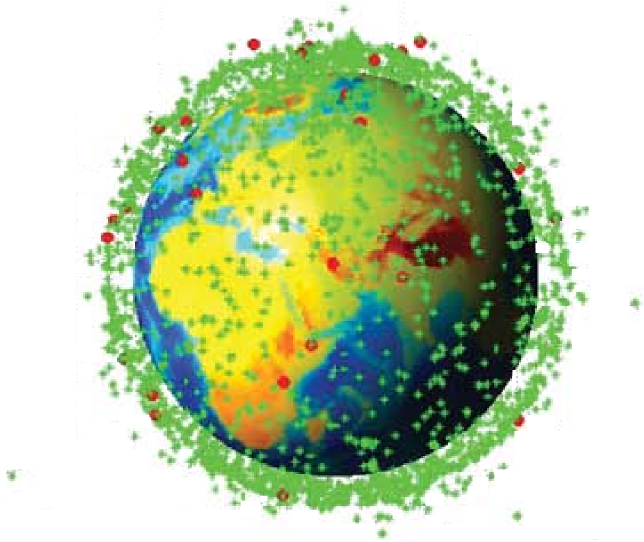

A snapshot of the environment surrounding the Earth shows the conjunctions that occurred in a 144-minute (less than two and a half hours!) time frame on 20 May 2009. Green dots mark the more than 2,500 conjunctions of less than 10 km, while the red dots mark the 31 conjunctions of less than 1 km.

satellite to avoid a collision with a piece of space debris. Although this is definitely a positive result (the satellite survives), the procedure leaves the satellite with less fuel to make future orbit corrections, potentially reducing its active lifetime. In fact, shifting position only extends the satellite's safety margin for about six hours before another conjunction of less than 10 km will occur, according to Lue.

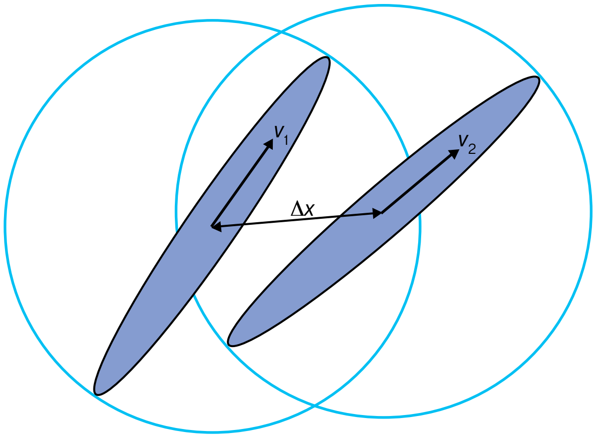

Although spherical conjunction, shown as light blue circles, might indicate that one or both of these objects should alter its path to avoid collision, the ellipsoids show that no corrections are necessary for this pass of these objects. In the figure, v_1 and v_2 are the (three-dimensional) velocities of the two objects, and Δx is the current separation.

For Lue, close is 10 km. If you set a sphere of 10 km radius around each object in space, these spheres will overlap with other objects every six hours or so. Extending this safety margin is the main thrust of Lue's work. Can all the objects that might cause damage be identified and characterized, and how accurately can their paths along their individual orbits be described? Newer and more sensitive satellite-tracking telescopes are helping to solve the first problem of locating objects down to less than 10 cm in size. Once each object is located, it needs to be continuously tracked to define its orbit. Now, Lue's analysis comes into play. The six-hour frequency conjunction mentioned above is for a simple sphere. Lue proposes that if the orbits can be defined more accurately, elongated ellipsoids of potential future locations of objects will not overlap as often and the ”ellipsoidal time between conjunctions” can be extended to 15 days—a significant improvement over six hours. With these tools—more accurate measurements of orbiting objects and improved algorithms for defining future locations of the objects—only those very close conjunctions will require a notification to satellite owners to suggest that they move the satellite to avoid the collision. As an added bonus to the collision avoidance, “there shouldn't be as many additional influxes of debris into the Space Catalog, “ Lue says.

Distance of closest approach	Number of spherical conjunctions per day	Number of ellipsoidal conjunctions per day	Time between spherical conjunctions per object	Time between ellipsoidal conjunctions per object
<100 m	3	3	5.9 years	5.9 years
<1 km	274	247	24 days	24 days
<10 km	27,271	547	6 hours	14.5 days

Columns 2 and 3 represent the total conjunctions in the current catalog over the course of one day for the specified distances of closest approach. For example, an object will approach another object within 1 km an average of 274 times per day. Columns 4 and 5 are the average times between conjunctions per object. Columns 3 and 5 are calculated for ellipsoids with 1-by-1-by-10 dimensions with the long direction listed in the table.

Lue has one final suggestion. “Don't put any more nonessential junk up there until we clean out some of the stuff that is already there.” Backing off from the Kessler limit requires cleaning up what already exists in orbit and creating a policy of minimal invasive actions when new satellites are deployed.

To envision the potential damage that a piece of debris could cause in a direct hit, consider the kinetic energy (potential to do damage on impact) of a car just before it hits the wall in one of the by now familiar auto safety commercials on television. The kinetic energy of a cube of aluminum 10 cm on a side moving at even 5 km/s relative to a satellite is 25 times greater than that of the car. The principal advantage for the satellite is that the cross section of debris/satellite impact is considerably smaller than that of the car and the wall.

1 D.J. Kessler and B.G. Cour-Palais, “Collision Frequency of Artificial Satellites: The Creation of a Debris Belt,” *Journal of Geophysical Research,* vol. 83, no. A6, 1978, pp. 2637–2646.

NANOTECHNOLOGY

Nanowire Single-Photon Detector Arrays

Detect a billion photons per second with low noise and high efficiency

Requirements for faster, high-sensitivity optical communication receivers have driven research efforts aimed at making higher-speed, more efficient photon-counting detectors. Superconducting nanowire single-photon detector technology has been developed for this reason through a collaboration between Lincoln Laboratory's Optical Communications Technology Group and MIT's Research Laboratory of Electronics.

Since the first demonstration of individual superconducting nanowire photodetectors in 2001, improvements have enabled single-photon detection that simultaneously achieves high detection efficiency, high speed, and low noise, resulting in a level of performance that has permitted a number of record-breaking experiments in high-sensitivity optical communication. This new level of performance was in large part enabled by the invention (by Dr. Eric Dauler, Dr. Andrew Kerman, and colleagues at MIT) of the interleaved, subwavelength-separated, nanowire single-photon detector array. Dauler states that these detectors are “uniquely suited to very efficient photon detection and counting of photons very precisely in time.”

The dimensions of a single wire segment are much smaller than the wavelength of light that must be detected, so even a diffraction-limited beam cannot be focused onto a very short wire. Thus, a nanowire is typically patterned to cover a larger area. Furthermore, the density with which wires are packed in this structure affects the absorptance of the detector, with a more tightly packed structure increasing the absorptance. Consequently, individual superconducting nanowire photodetectors typically have small active areas. They also cannot resolve the position, wavelength, or number of photons in an optical pulse, and require a minimum of several nanoseconds between detection events. Superconducting nanowire photodetector arrays, with their subwavelength gaps, address these shortcomings of individual nanowires and achieve a level of performance beyond that possible with a standard array configuration. Superconducting nanowire photodetector arrays also eliminate the need for complex optical coupling that is commonly necessary to increase the fill factor for other photon-counting array technologies.

Although a trade-off clearly exists in choosing nanowire length for a single detector (speed and yield improve for shorter nanowires while longer nanowires offer a larger coverage area for multimode signals), the superconducting nanowire photodetector array simultaneously combines many of the advantages of short and long nanowires. One realization of a superconducting nanowire photodetector array utilizes a spatially interleaved array of four serpentine, superconducting nanowires that occupy a region 7 to 20 μm in diameter to form a single optical active area. In this case, each nanowire is four times shorter than would be needed if only a single nanowire had been used to cover the same total area. Therefore,

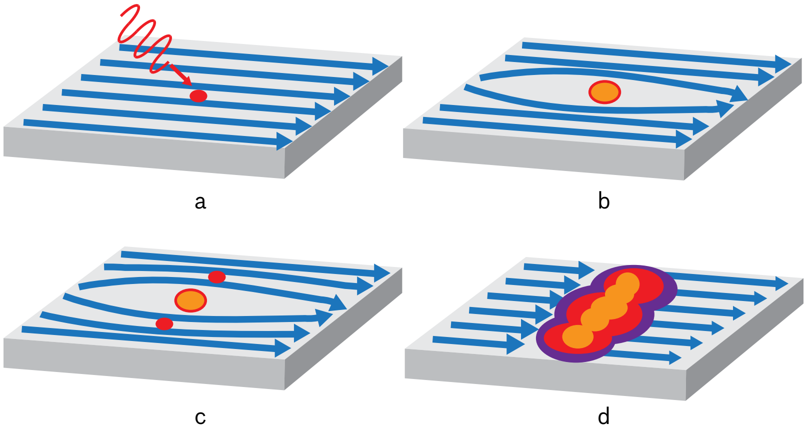

A photon is incident upon a nanowire (a) biased with a direct current near the critical current value, heating the nanowire and forming a localized hot spot (b) where the superconductivity is disrupted. The supercurrent diverts around this spot (c). For a sufficiently narrow wire (<100 nm), the local current density on either side of the hot spot exceeds the critical current density, causing a resistive region to span the entire cross section of the nanowire (d) and inducing a measurable voltage across the device.

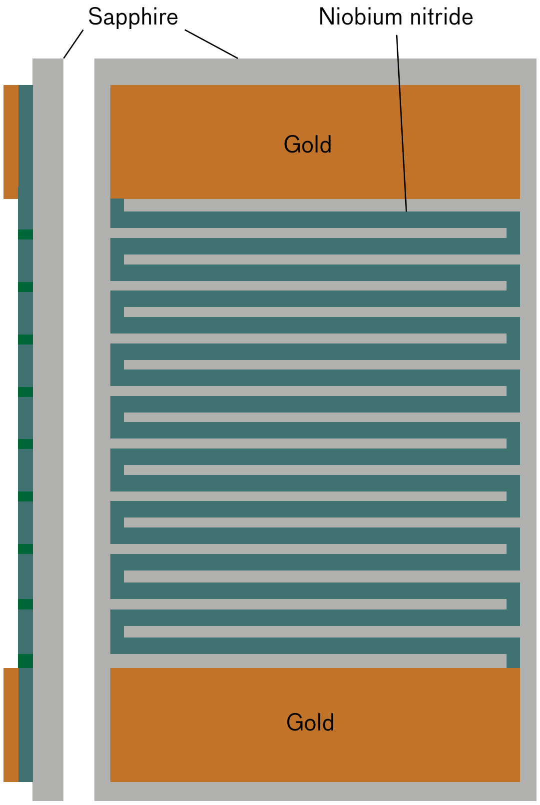

A superconducting nanowire photodetector can be wrapped to cover a larger region of detection, but longer wires require greater recovery, or reset, time before detecting additional photons.

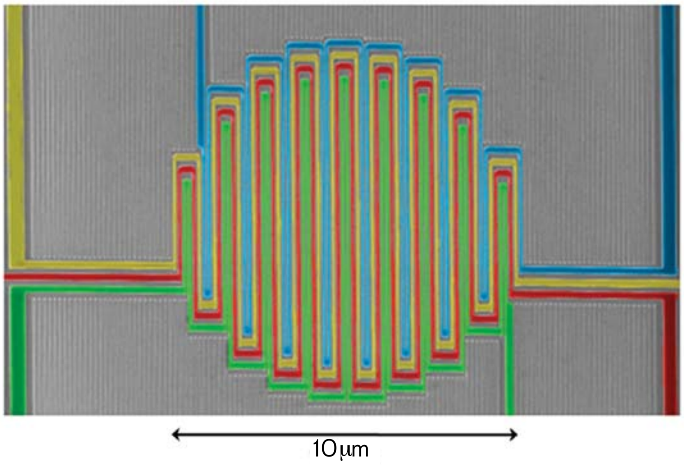

Each nanowire in a superconducting nanowire photodetector array (shown here in a false-color scanning electron microscopy image with four individual, spatially interleaved wires) can recover more quickly following a photon detection, and the other wires can be immediately available to detect subsequent photons.

each of the four shorter nanowires is faster and can be produced with higher yield, while the combination of elements covers a large area with tightly packed wires.

When illuminated by a single photon, only one of the interleaved nanowires registers a detection event. This single photon creates a hot spot that results in a large change in resistance, producing a signal that dominates all noise sources. The nanowire then quickly cools back down and the superconducting current is able to begin flowing in the nanowire again. The characteristic inductive time for the current in the detector to recover to its initial value depends on the length of the nanowire, and determines the reset time after a detection event before the next photon can be detected. However, during this time, the remaining nanowires are available to detect the next photon. This spatially multiplexed approach allows detection of photons at a higher rate than a single, superconducting nanowire would support, both because the individual nanowires are faster and because the array elements can count in parallel. Additionally, using multiple nanowires provides information about the number of photons in an optical pulse (up to the number of array elements) and permits the detection of photons separated by time periods much less than the reset time of a single nanowire.

One of the leading applications of this technology is in ultrahigh-sensitivity optical communication. The existing implementation of the superconducting nanowire photodetector arrays will permit photon-counting communication at faster rates than were previously possible. Whereas other existing near-infrared photon-counting detectors in the optical communication bands provide either high speed (e.g., photomultipliers) or high sensitivity (e.g., infrared avalanche photodiodes) alone, the Lincoln Laboratory superconducting nanowire photodetector arrays combine excellent performance in both sensitivity and speed.

The Lunar Laser Communications Demonstration rate of less than 6×10^8 photons per second is still not pushing the detector to its limit.

“We’re focused on really long-distance communication—lunar and planetary distances,” Kerman explains. The superconducting nanowire photodetector technology

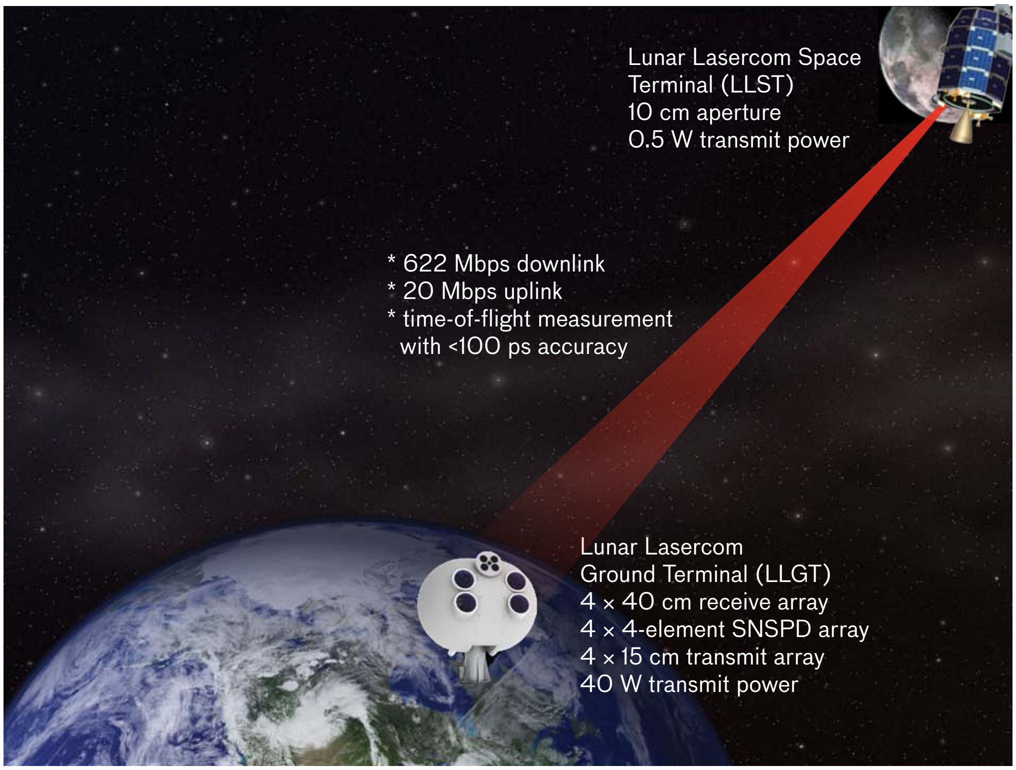

The Lunar Laser Communications Demonstration will demonstrate a long-distance, photon-counting, laser communication system capable of 622 Mbps optical downlink. The superconducting nanowire single-photon detector (SNSPD) arrays will be used in the Lincoln Laboratory ground terminal.

will be used by Lincoln Laboratory for the NASA-funded Lunar Laser Communications Demonstration (LLCD). “The telecom industry isn’t interested in such sensitive devices since they can send signals from amplifier to amplifier along the chain when they need it. We don’t have the ability to put in an amplifier between here and the moon or Mars.”

Past successful NASA space exploration missions have relied on radio-communication links to interplanetary spacecraft as well as to Earth-orbiting satellites. However, this modality will not support future missions that will deploy more sophisticated instruments (e.g., synthetic aperture radars) requiring data rates of much greater capacity. Optical networks represent the future of deep-space communications and will make it possible to collect ambitious data products from more distant destinations. Highly sensitive single-photon detection is a key enabler for this next generation of optical communication technologies.

Dr. Bryan Robinson, the lead system engineer on the LLCD program and a colleague of Dauler and Kerman, says that a probe used in a recent mission to Pluto with a 600 bit/s radio-frequency link will require approximately nine months to download the data collected during a few-week encounter with Pluto. Dauler responds, “A higher data rate communication system would enable more data collection. A mission’s value and flexibility could be increased by near-real-time downloading of information.”

The LLCD—part of the moon-orbiting Lunar Atmosphere and Dust Environment Explorer satellite experiment scheduled for launch in 2013—will demonstrate a long-distance, photon-counting, laser communication system capable of sending information at 622 Mbits/s from the moon to Earth. The superconducting nanowire photodetector arrays, located in a transportable ground terminal, will be key elements in this communication link. “The photon detection rate for this link is several hundred million photons per second,” Dauler says. “This is still not pushing the detector to its limit.”

The superconducting nanowire photodetector array has numerous intrinsic advantages—high performance in both sensitivity and speed, the largest photon-counting rates available, relative ease of fabrication, simple digital post-processing to obtain photon-number information, and precise photon timing information. These advantages make the technology adaptable to other potential applications:

- Single-photon-counting detectors for quantum key distribution systems demonstrated at 1.85 Mbits over 100 km distances—rates 100 times faster (for a fixed fiber-optical cable length) than previously reported
- Ultrasensitive, time-resolved spectroscopy. Biological systems often exhibit fluorescence, which is sometimes in the infrared spectrum. For example, the singlet O_2 fluorescence at 1.27 µm wavelength is important in cellular metabolic processes. Existing biological imaging systems focus on the visible domain because of the difficulty of working with long-wavelength imaging arrays. The availability of a high-speed, sensitive, infrared detector may enable new fundamental studies.
- Basic science experiments in quantum mechanics, including quantum imaging, quantum measurement, and quantum entanglement
- A high-resolution, noncontact method to detect flaws in high-speed, very-large-scale integrated complementary metal-oxide semiconductor circuitry
- Imaging technology. High-speed video imaging could be achieved with an array of nanowire photodetectors spread over the focal plane of an optical camera system

AIRPORT SAFETY

Runway Status Lights

Protecting aircraft when they are most vulnerable—during takeoff and landing

Jim Eggert and Eric Shank of the Surveillance Systems Group are happy to have assisted in the development and deployment of a system that aids in preventing runway incursions at several airports. Much work has focused on improving safety during flight, but at the point where aircraft come closest at high relative velocities—on takeoff and landing—more can be done to improve the safety of aircraft and passengers. Runway incursions, when an aircraft or vehicle is on a runway without permission from air traffic control, are a daily occurrence in the United States. Preventing runway incursions that lead to accidents has been on the National Transportation Safety Board's "Most Wanted List" for over two decades.

Although the domestic aviation system in the United States is one of the safest man-made transportation systems devised, accidents still occur. One of the aviation industry's most urgent safety concerns is that of high-speed collisions involving aircraft on runways. ("There are fender-benders [low speed interactions], but they usually don't endanger the flying public," according to Shank.) Such collisions may be between two aircraft or one aircraft and a service vehicle.

Takeoffs and landings at airports are critical points in an aircraft's flight. An airplane accelerates down a runway on takeoff, trying to gain enough speed to achieve lift off, and it would be very difficult to stop it if something appeared on the runway in front of it. Similarly, landing aircraft are slowing down, trying to reduce their lift, and at some point, they cannot accelerate again to regain altitude. In either case, an object on the active runway would certainly be problematic. Eggert and Shank and their associates are working with the Federal

Takeoff Hold Lights indicate that there is some obstruction (in this case crossing traffic), so the pilot should hold position until the runway is cleared.

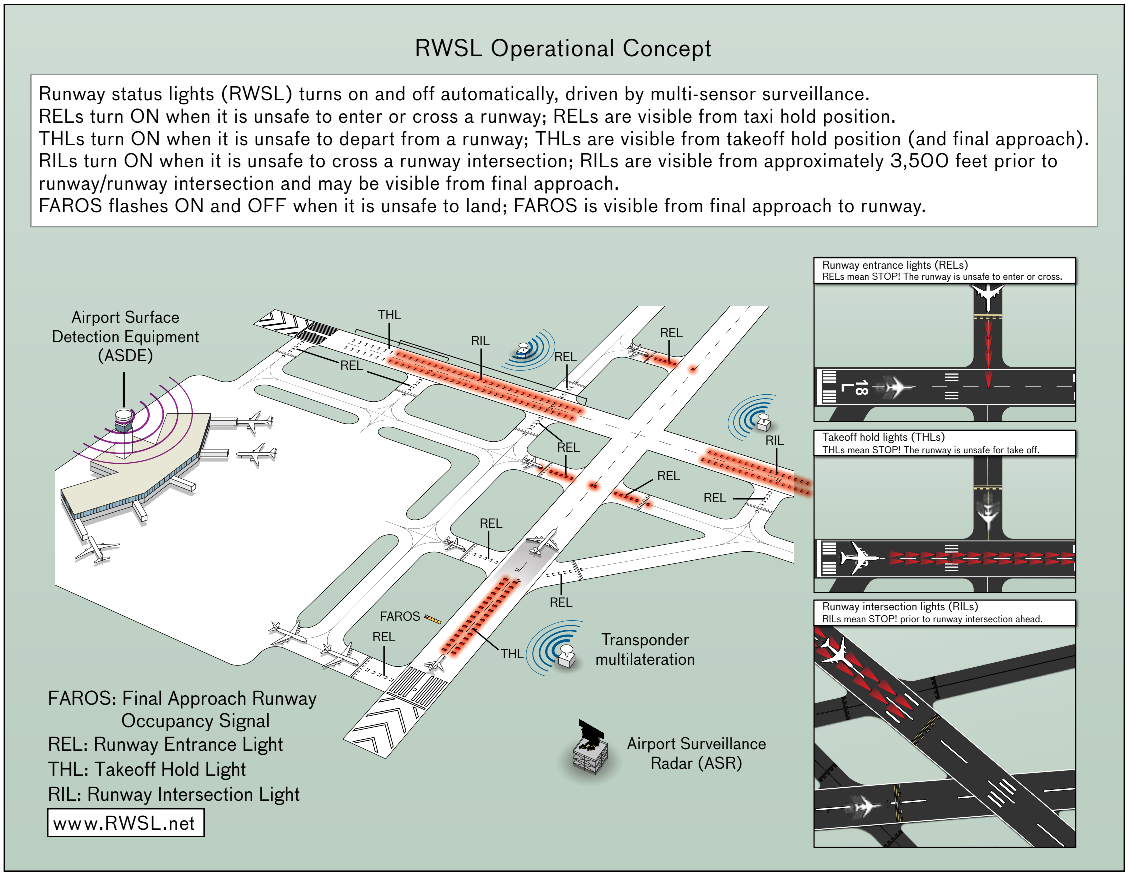

RWSL Operational Concept

Runway status lights (RWSL) turns on and off automatically, driven by multi-sensor surveillance.
RELs turn ON when it is unsafe to enter or cross a runway; RELs are visible from taxi hold position.
THLs turn ON when it is unsafe to depart from a runway; THLs are visible from takeoff hold position (and final approach).
RILs turn ON when it is unsafe to cross a runway intersection; RILs are visible from approximately 3,500 feet prior to runway/runway intersection and may be visible from final approach.
FAROS flashes ON and OFF when it is unsafe to land; FAROS is visible from final approach to runway.

Runway entrance lights (RELs)
RELs mean STOP! The runway is unsafe to enter or cross.

Takeoff hold lights (THLs)
THLs mean STOP! The runway is unsafe for take off.

Runway intersection lights (RILs)
RILs mean STOP! prior to runway intersection ahead.

FAROS: Final Approach Runway Occupancy Signal
REL: Runway Entrance Light
THL: Takeoff Hold Light
RIL: Runway Intersection Light

www.RWSL.net

A typical airport environment supplied with an RWSL system will have four types of status indicators. The REL, THL, RIL, and FAROS components of RWSL are shown in their relative locations on and near active runways.

Aviation Administration (FAA) to eliminate runway incursions during takeoffs and landings.

Runway accidents often develop so quickly that mediation by air traffic control (ATC) personnel is often impossible or ineffective. Timely communication of runway-safety information directly to pilots is often required to avoid a runway incursion or collision. The concept of notification relies on the ability to alert at least one of the aircraft or vehicles in the conflicting scenario. In some cases, for increased safety, redundant indications are provided to everyone involved.

Currently, the most effective direct notification system at towered airports is the runway status lights (RWSL) system, developed for the FAA by Lincoln Laboratory. RWSL indicates to a pilot when a runway is unsafe by turning on special red lights embedded in the pavement in full view of the pilot and other nearby personnel, or by flashing lights to pilots on approach to the airport. The RWSL system operates independently from the clearances issued by ATC and thus serves as an independent layer of safety. RWSL meets a long-standing, well-defined safety need to help prevent runway incursions and accidents by combining current ground-based radar and multilateration technology with advanced processing to control in-pavement lights that directly alert pilots to runway collision hazards. Of the technologies specifically addressing runway incursions, RWSL provides the most timely, most effective, and most highly automated technology to notify pilots and vehicle operators on the airport surface of potential incursions.

Eggert recalls that their initial work was a pair of prototype simulations at Boston's Logan International Airport. The first one involved simulated surveillance, pseudopilots (a computer display that allowed a technician to control several simulated aircraft), and a controller. This simulation allowed the concept to be tested with realistic controller-pilot communications and aircraft motions.

The second prototype used real, live surveillance data and showed runway status light operation as it would occur, but only on a computer display and on a model board with fiber-optic lights, not with real lights on the airfield.

The second simulation involved real data, a simulated controller, "pseudopilots," and a model board with light-emitting diodes. "We were given a room in the tower so we could watch the same data that the controllers saw," Eggert says. The conclusion of the initial work was that the radar technology at the time was insufficient to maintain the necessary high degree of proper signalling with a minimum of false alarms (which would reduce runway capacity). "If something is happening that would make it dangerous to continue what they are going to do, we want them to know," Shank says. "Otherwise, we don't want to interfere with operations because delays cost money." However, as technology improved, RWSL were installed at several airports. The first fully functioning operational prototype was installed at Dallas/Fort Worth International Airport (DFW), Texas, where a prototype Airport Surface Detection Equipment (ASDE-X) radar was installed. "It reduced runway incursions by 70% when tested at DFW," according to Eggert.

The four components of RWSL (as shown in the sketch of a typical airport) are Runway Entrance Lights (RELs), Takeoff Hold Lights (THLs), Runway Intersection Lights (RILs), and Final Approach Runway Occupancy Signal (FAROS). Each of the first three types of light has only two states: "On" (lights illuminate red) and "Off" (lights not illuminated). No third state exists; RWSL never, for example, displays green lights. The fourth component, FAROS, also has two states: "On" (lights illuminated white over red) and "Flashing" (lights flashing on and off).

- RELs are placed at runway/taxiway intersections and are visible to a pilot taxiing toward a runway. They indicate to the pilot if it is unsafe to enter or cross a runway because it is currently or will soon be occupied by high-speed traffic, such as an aircraft taking off or landing, and the pilot should stop immediately. (For simplicity of description, the discussion in this note emphasizes aircraft-to-aircraft encounters, but it should be kept in mind that RWSL has also been shown to be effective in averting aircraft collisions with surface vehicles.)
- THLs are placed on the runways to be visible to the pilot in position for takeoff. They indicate to the pilot that it is unsafe to take off because the runway ahead is occupied by another aircraft. If a pilot is holding on a runway when THLs illuminate red, the aircraft should remain in position. If a takeoff roll has begun when a pilot observes illuminated THLs, the pilot should stop the aircraft and notify ATC that the plane has stopped because of red THLs.
- RILs are placed on runways approaching an intersection with another runway to indicate to a pilot in a takeoff or landing roll that the intersection ahead is unsafe to enter or cross because there is a potential conflict at the intersection. When RILs illuminate red, the pilot or vehicle operator should stop before the intersecting runway.
- FAROS is a flashing signal imposed on the already-existing precision approach path indicator (PAPI) lights, visible to aircraft on final approach to a runway. They indicate to pilots that the runway is occupied, and the pilot should visually acquire the other traffic and may have to contact the tower to verify clearance to land or, absent that verification, go around instead of land.

RWSL works seamlessly with existing and planned ATC procedures. The RWSL system is effective because an indication of a conflict is

- Transmitted directly to the pilot(s) involved.
- Generated by computer logic.
- Is not dependent on the visible detection by controllers and/or vehicle or aircraft crew members to enhance safety during night operations or periods of restricted visibility.
- Does not depend on the availability of clear audio channels.

> It's easily understood. Red means stop! People understand that.

At the request of the FAA, Lincoln Laboratory reviewed runway incursions in the United States between 1997 and 2000 at 100 of the busiest airports, concentrating on those incursions that involved at least one large passenger jet and were classified as "high hazard" or had a miss distance less than 100 feet. The study determined that RWSL might have prevented or mitigated 75% of the 167 identified incursions. The study suggested that the efficacy of RWSL stems from their ability to directly alert pilots of the runway status with minimal latency. Furthermore, RWSL helps prevent the occurrence of incursions—the predecessors of accidents—by increasing the situational awareness of pilots on runways and taxiways.

According to the FAA and the NTSB, RWSL is a viable and important technology for reducing runway incursions. In addition, RWSL has gained widespread support among user groups. It requires no human processing or warning, does not increase the ATC procedural workload, and does not interfere with other pilot procedures and tasks. Pilots, pilot union officials, air traffic management, and the airport operator at DFW all agreed that RWSL works as intended and has no known negative impact on capacity, communication, or safety. NTSB officials stated that RWSL is a promising technology for addressing its long-standing recommendation to provide pilots with direct warnings of potential runway conflicts.

As a result of successful operational evaluations of prototype RWSL systems at DFW and San Diego International airports, the FAA announced in 2007 its decision to install RWSL at 23 major airports in the U.S. National Airspace System. The FAA contracted with industry to produce the Lincoln Laboratory–certified system for delivery in 2009 to Los Angeles International Airport. The same system has also been deployed at Boston's Logan International Airport. Lincoln Laboratory is currently working with industry and the FAA to complete the technology transfer of RWSL.

"It has been a great project to work on," Eggert concludes. "It's easily understood. Red means stop. People understand that."

"We still have to be very efficient in getting planes into the air," Shank says. In case an error is made, all the pilots have to do when they see a stop red light is to contact the tower for further information or instructions. "RWSL makes our safe aviation system even safer," Eggert concludes.

PHOTON DETECTORS

Ultrasensitive Two-Dimensional Photodetection

What information can be associated with the detection of a single photon?

Where and, more importantly, when did you see that flash of light? Specifically, when did that single photon arrive at the detector? And how quickly can the detector recover to see the next one?

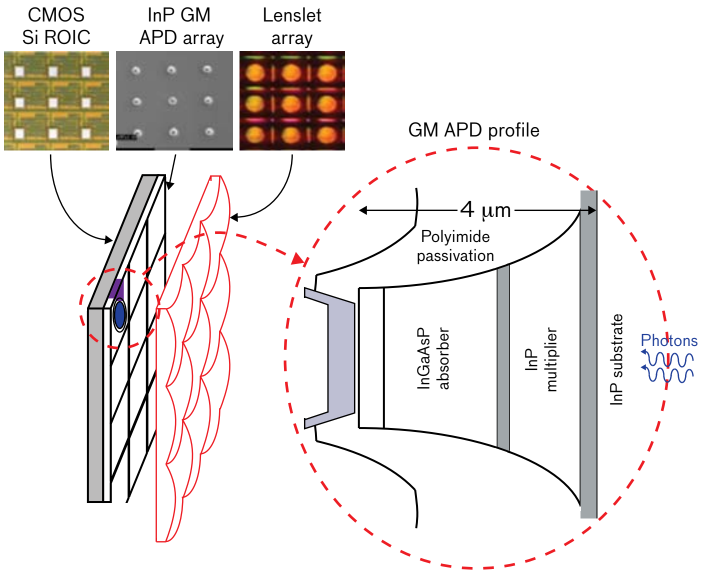

Each pixel of the array comprises a stack of components. Photons arriving from the right pass through an array of lenses and hit the GM APD array. Each pixel of the array has its own timing circuit directly behind it, so digital counting signals are generated at the pixel, thus allowing the pixel to be quickly reset. In this figure, CMOS stands for complementary metal-oxide semiconductor, ROIC is readout integrated circuit, and InP and InGaAsP are indium phosphide and indium gallium arsenide phosphide.

These questions and others are the bailiwick of Richard Marino and his associates in the Active Optical Systems Group. The solution for getting answers to these questions is a focal-plane-detector array comprising Geiger-mode avalanche photodiodes (GM APD FPA).

Back in 1991, Marino recalls, there was a need, expressed by the Department of Defense, for a very smart missile that could quickly and automatically distinguish between a true target and a decoy. The missile sensors had to be very capable, and yet small and light in order for the missile to be quick enough to engage the target. The sooner the missile could identify the target, the sooner it could divert its path to hit the target. The sensor system could use a laser to probe the targets and decoys, but the performance of the laser radar (or ladar) was crucial. The question came down to, "From how far away can the sensor identify the target?" and led to the fundamental question, "How much information can be associated with the fewest amount of photons detected?"

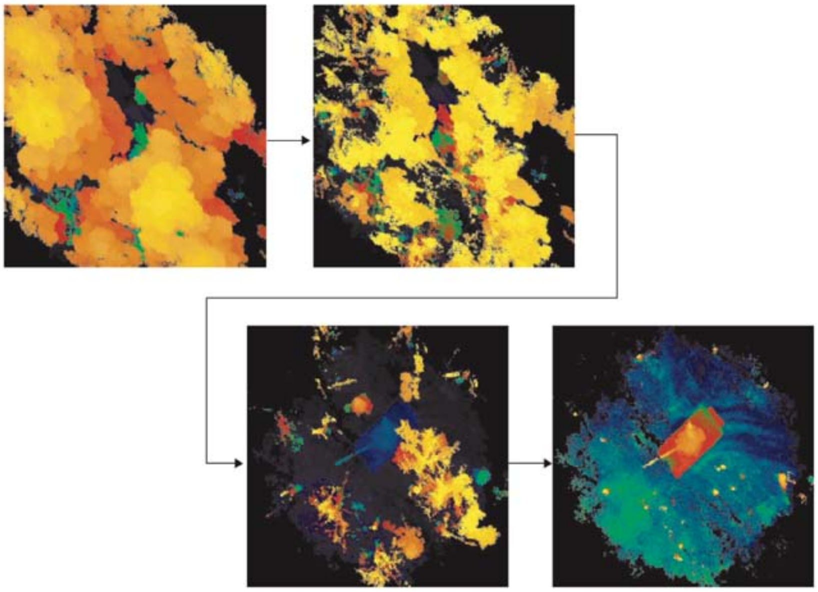

Foliage penetration is possible through multiple images from various angles. Even though an object is completely covered to the eye (e.g., tree foliage above a vehicle), laser radar signals can visualize any level (distance from the detector) of information. This sequence of images shows the effect of cropping the three-dimensional data by successively eliminating pixels above a certain height. The first image (upper left) contains treetops that are eliminated by the fourth image (lower right) to reveal a hidden object below the tree canopy.

The concept of a Geiger-mode photon-counting detector array that could measure the three-dimensional shape (resolved in angle, angle, and range) and orientation of the target and decoys was proposed. Such a sensor could reduce the requirements for size, weight, and power by maximizing the efficiency of the optical signal receiver. A compact, intelligent, integrated detector array was needed. "There's skepticism even today about the GM APD technology," Marino says. The nonlinearity of the GM APD and the fact that it isn't counting enough photons for traditionally "good" statistics are typical concerns in ladar sensor engineering. "Engineers like linear systems," Marino states. However, the technology behind GM APDs is almost perfectly nonlinear. "Timing is everything," Marino continues. The arrival time of a single photon at a detector determines the distance to the corresponding spot on the object.

Engineers like linear systems, but photons are discrete. We use GM APDS as nonlinear photon-to-digital converters.

Brian Aull and his associates in the Advanced Imaging Technology Group developed GM APD arrays that have the ability to detect and time-stamp single photons by using their unique, independent, digital time-of-flight counting circuits (at each pixel) and their extremely high internal detector gain. The independent time-of-flight measurement for each pixel has a timing quantization of 500 picoseconds (equivalent to a 2 GHz effective clock rate). Using the detector in this binary response mode, where time of arrival is more important than signal intensity, simplifies the signal processing by eliminating the need for analog-to-digital converters and corrections to varying responses to input intensity. Simplifying the detection process in this way, and maximizing single-photon-detection efficiency helps in reducing the required size, weight, and power requirements. Marino is

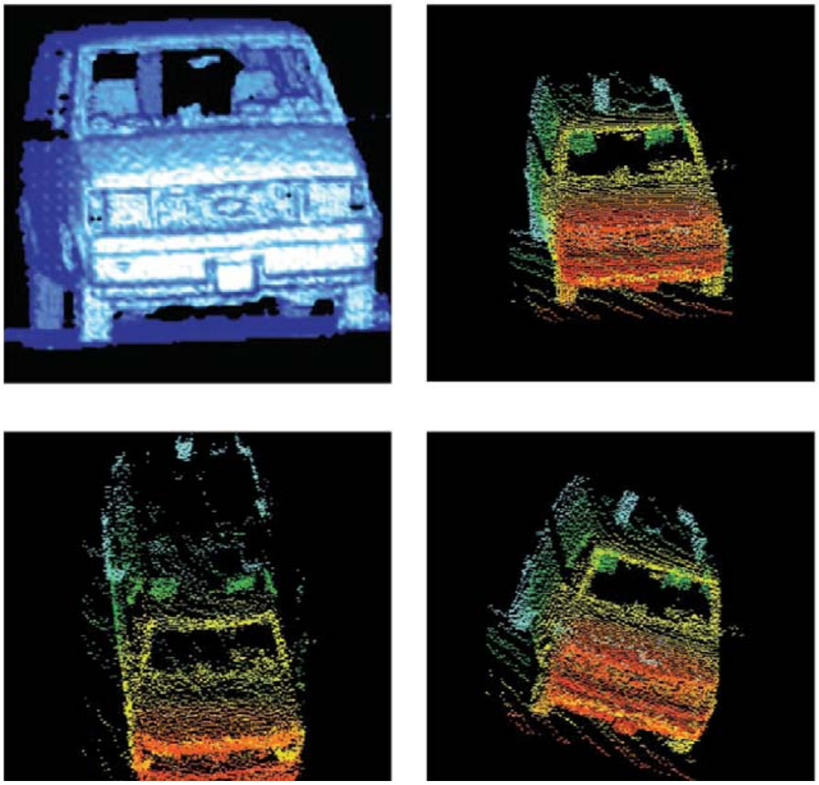

These images of a Chevrolet van were obtained from a prototype three-dimensional laser sensor. In the upper left is a three-dimensional model rendered from the angle-angle-range data. The other three renditions are point clouds viewed from different aspects. Rotation of the color-coded image better reveals shapes, sizes, and relative positions of different parts of the van.

also concerned with possible false alarms with such sensitive devices. In order to reduce unwanted detections from random background light, GM APDs are not held in their wait state indefinitely—they are turned on only during the expected detection time (range to objects).

A single GM APD can be thought of as a photon-to-digital converter that produces a digital-logic-compatible voltage transition in response to a single incident photon. In this way, GM APDs completely eliminate many of the traditional types of noise (e.g., read noise, amplification noise) involved in photon detection with analog receivers. In the GM APD arrays developed at Lincoln Laboratory, each pixel is mated to a digital CMOS (complementary metal-oxide semiconductor) timing circuit that measures the arrival time of the photons.

The independent time-of-flight measurement for each pixel has a timing quantization of 500 picoseconds (equivalent to a 2 GHz effective clock rate). One primary application of the GM APD imager is as a detector array in a three-dimensional imaging laser radar (ladar) camera. A ladar camera uses a very-short-pulse laser (a typical laser pulse width is 1 ns) to illuminate an object and a GM APD optical receiver to simultaneously image the reflected light and measure the time of flight of each photon for each pixel in the image. The resulting three-dimensional data (x- and y-coordinates corresponding to the pixel position in the array and a z-coordinate corresponding to its range) can be mapped, with a color spectrum representing the range in the image to produce a three-dimensional “point cloud” depiction of the object. Such imagery is useful for looking behind partially opaque material (e.g., foliage) as well as for target identification and feature extraction.

The second innovation that made photon-counting ladar sensors so successful was the microchip laser. These lasers proved effective in reducing the camera size and weight while still providing “enough energy to get the information you want,” according to Marino. Under normal circumstances, to produce sufficient statistics, you need high intensity or you need to integrate over long time periods if the intensity is low. With the current microchip lasers, even a relatively small transportable or space-qualified camera has plenty of intensity for photon detection. Flash ladar, a single big pulse, certainly has sufficient intensity to acquire adequate data, but these lasers tend to have low pulsing rates (10 to 100 pulses/s). If the sensor receiver requires a strong signal to make a detection, then there will be noise from speckle interference in the intensity data, and multiple pulses are usually averaged to reduce the effects of speckle, according to Marino. “Instead, we typically use lasers that pulse at a very high rate (10,000 pulses/s or greater) and operate with less than one photon per pixel per pulse on average.”

"Signal-to-noise ratio (SNR) is commonly used to determine the performance or quality of a measurement, but SNR usually refers to intensity," Marino says. The SNR figure of merit isn't an obvious concept in the GM APD data. "With a GM APD, a photon is either detected or not—a zero or a one." Marino does consider that the error in range is a measurable value that can be applied as a figure of merit for the GM APD.

With a GM APD, a photon is either detected or not—a zero or a one.

One application that has been successfully demonstrated and utilized is an airborne three-dimensional ladar camera. This device collects multiple three-dimensional data sets from a collection of viewing angles and then registers and combines the images, as in a jigsaw puzzle. By adjusting the "threshold" of the photons' return times, the color coding of the images can be selected to remove the obscuring materials (e.g., foliage) from the image to reveal the objects of interest beneath. An illustration of this process uses data obtained from a helicopter. Multiple three-dimensional images are collected from different viewing angles, spatially registered to form a three-dimensional point cloud, and displayed with color representing relative height.

The GM APD array cameras have been used for imaging wide areas and urban environments, as well as for detecting change within an image. Future applications are wide open. With these cameras, it is possible to rapidly capture true three-dimensional data of an entire construction area, perform accurate land surveying, and create precise three-dimensional surface models of solid objects. This sensor technology in its current form could be used for detection and tracking of moving objects for border patrol, for robotic vision, and for navigation of fully autonomous air, land, sea, and underwater vehicles.

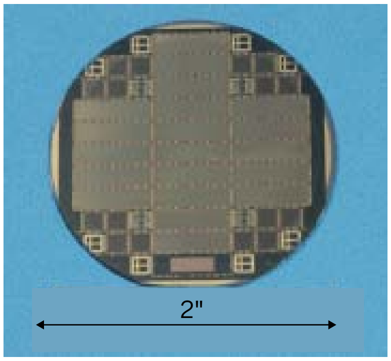

The wafer above contains twenty 256 × 64 FPA chips.