Early Warning Chemical Sensing

Juliette A. Seeley and Jonathan M. Richardson

The threat of chemical weapon attacks has prompted the defense community's development of standoff chemical sensors, designed to provide advance warning of an attack from a distance. Lincoln Laboratory has assessed the mission requirements for two relevant applications—wide-area chemical surveillance and inexpensive fixed-site protection—and developed sensor concepts optimized for each mission.

» **The nation faces a growing concern that** terrorists or state actors will resort to the use of chemical or biological warfare agents against military or civilian targets. The primary defensive action available for this threat is to detect the deadly agents before they contact their intended victims. Given advance warning, people can evacuate a contaminated area, administer prophylactic medication, or don protective equipment. Moreover, timely awareness of an attack will make it more likely that victims who are exposed will receive appropriate treatment.

Chemical agents may be delivered either as a vapor or aerosol or in combination, depending on their physical properties, and various agents will have different persistence following an attack. Military countermeasures may be developed for fixed sites or to be deployed during tactical operations. Normally, tactical deployment implies protecting a moving force, which has proven to be more difficult than protecting a fixed site. Attacks may be either overt, using chemical munitions, or covert, using a variety of improvised methods. Attacks may also be either exterior or interior, with exterior attacks classified as either on target or off target (localized or from a point upwind). When the chemical target is a military installation or other outdoor target, a variety of atmospheric and other relevant considerations also come into play, including the type of location (urban, rural, desert, or maritime), time of year, and local topography.

Chemical sensors can be categorized as either "point," meaning that they sense local conditions, or "standoff," meaning that they sense threats at a distance. While point sensors have been demonstrated to have very high sensitivities, they have a significant drawback: protection over a wide area requires deployment of a large number of individual sensors. Standoff systems make it more practical to monitor a wide area with fewer sensors and make it possible to detect chemical releases from remote locations such as aircraft.

Lincoln Laboratory has developed two long-wave infrared (LWIR) standoff chemical-sensor concepts, optimized for two particular missions: one for wide-area chemical surveillance, the other for inexpensive protection of fixed sites—particularly civilian facilities such as airports, arenas, and large office buildings. The Laboratory's approach has been to first address the logistical constraints of each mission and then trade sensor parameters so as to optimize performance within the given logistical constraints. We describe Lincoln Laboratory's development process for each sensor from conceptual design to its system test.

LWIR Spectroscopy

The predominant technique used for standoff chemical sensing—LWIR absorption spectroscopy—relies on the distinctive absorption signatures that most chemicals exhibit in the spectral region between 8 and 14 micrometers [1]. Typically, larger molecules (such as some chemical warfare agents) will have a few broad spectral features. By contrast, the lighter molecules found in many toxic industrial chemicals (TICs) have numerous sharp features that are more difficult to detect. (However, some TICs, notably chlorine gas, have no IR absorption features at all and thus must be detected by other means.) IR absorption spectrometry is best suited for detecting vapors but can also sense chemical aerosols. Recent work has shown that IR absorption spectroscopy can even detect bioaerosols if the agent is present in sufficient quantity [2].

Spatial information obtained with the use of a hyperspectral imager can add more certainty to detection statistics, potentially reducing the incidence of false alarms.

IR absorption spectroscopy is a powerful technique for determining the quantity and identity of many molecular species; it has been used for more than 50 years in numerous laboratory and field instruments. In a typical laboratory IR absorption experiment, an IR source illuminates the sample, which selectively absorbs wavelengths of light corresponding to molecular vibrational transition energies. The spectrometer separates the transmitted radiation into its component spectral bands, which are then sensed by the detector element or elements. The spectral features of the recorded data indicate the identity of the sample.

A LWIR system can be used as a point sensor, detecting chemicals only at one particular location. Alternatively, we can use the LWIR system in standoff mode to monitor for the presence of chemical agents from a safe distance [3]. Detection in the LWIR region is conducive to remote detection because of a transmission window in the LWIR that allows for light in this region to be transmitted by the atmosphere. A standoff system increases the area that is protected by an IR sensor in a given amount of time—that is, the area coverage rate (ACR).

Standoff LWIR sensing can be either active or passive. Active sensing relies on an IR source such as a hot filament, a laser, or another emitter. One option is to separate the source from the sensor so that the sensor is sensitive to the line of sight between the source and sensor. Such LWIR line-of-sight sensors have been used in the past for environmental monitoring [4, 5], though such applications are often concerned with trace quantity analysis, demanding detectors with the highest achievable sensitivity. Alternatively, the source and sensor can be collocated, and the radiation scattered from the plume back to the sensor or reflected from a hard background or other reflector located behind the plume. The active sensor under development at Lincoln Laboratory is a line-of-sight sensor that uses a hot filament source.

Passive sensors, by contrast, acquire signal from the blackbody radiation of the scene (land, sea, sky, or combination); at ambient temperature, significant blackbody radiation in the 8 to 14 μm region is available for passive detection. A familiar analogy is that of a visible camera without a flash, whereas an active sensor is analogous to a camera with a flash. LWIR radiation emitted by the background is absorbed and emitted by all intervening materials, including a threat cloud, if present, as well as other ambient vapors or aerosols within the field of view. The differential signal contributed by a threat cloud

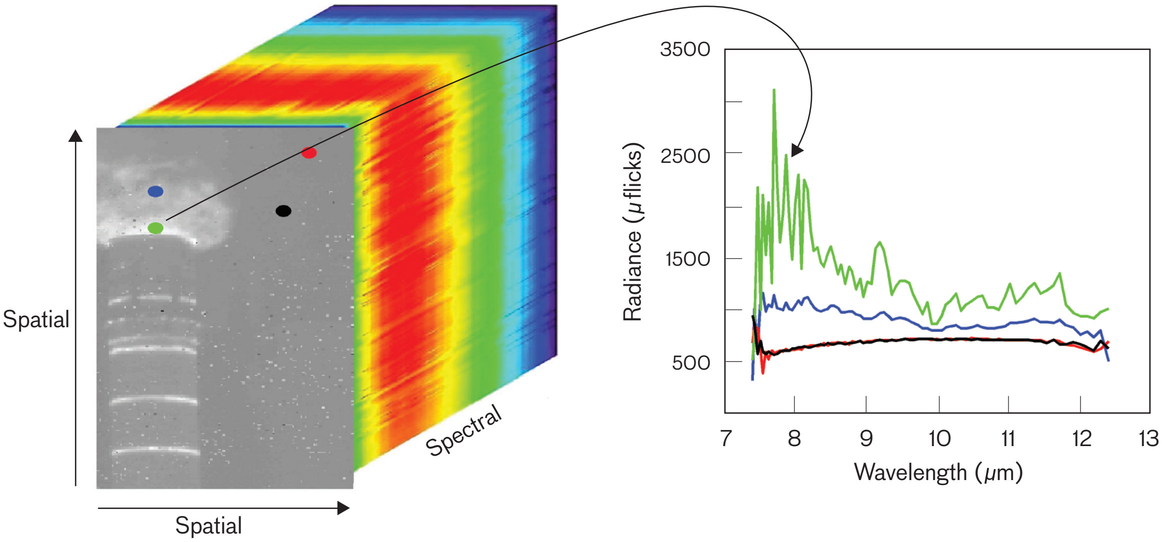

FIGURE 1. An imaging spectrometer provides simultaneous spatial and spectral information about a scene. Its output, a hyperspectral data cube, contains spectral information at every spatial point as well as an image at every spectral band. The figure on the left depicts one spatial slice of a cube at one particular spectral band. Color-coded spectra of specific pixels in the image are shown on the right. The green pixel is particularly warm and exhibits the spectral structure of the gas emitted from the stack. The red and black pixels, corresponding to cold sky, have lower radiance values.

will depend upon the relative temperatures of the various absorbers/emitters, their vibrational modes, and the concentration and length of the threat cloud along the sensor line of sight.

Various advantages and disadvantages are associated with either active or passive detection. Since passive systems acquire ambient radiation from the scene instead of being limited to a line of sight determined by the location of the source and detector, they can view a larger area at any given time, resulting in a larger ACR. Active sources, on the other hand, are more stable than the fluctuating scene radiation. The use of an active source can also provide a higher signal, making the use of less sensitive, low-cost detectors a viable option.

The military has been developing passive LWIR sensors for standoff chemical sensing since the late 1950s. The most recent version is the Joint Service Lightweight Standoff Chemical Agent Detector (JSLSCAD) [6], which is based on its predecessor, the M21 Remote Sensing Chemical Agent Alarm (RSCAAL). JSLSCAD is a fully automated, multi-mission, single-pixel, passive IR system capable of discriminating between chemical agents and other species in a complex battlefield environment. While the performance of JSLSCAD is adequate for certain applications, there are many advantages provided by a multi-pixel sensor, as described below.

One way to enhance ACR, while maintaining high spatial resolution, is to use a detector with multiple pixels—that is, a focal-plane array (FPA). With appropriate imaging optics, this type of sensor functions as a hyperspectral imager, providing spectral information at each pixel location. Such a hyperspectral imager generates what amounts to a three-dimensional data cube, with two spatial dimensions and one spectral dimension, as depicted in Figure 1. In addition to the benefit of improved ACR, spatial information obtained with the use of a hyperspectral imager can add more certainty to detection statistics, potentially reducing false alarms. For example, spatial filtering may be used to exclude pixels that appear spectrally similar to a threat agent but are not consistent with the physical shape of a plume. It is typically true that isolated pixel detections are a result of noise. Several hyperspectral instruments are under development for use in plume detection [7, 8], as well as other surveillance applications in the visible region [9]. For multi-pixel LWIR hyperspectral imaging to be adopted for military use it will have to be rugged as well as affordable.

Radiative Transfer

To successfully detect a chemical plume, the system designer must be able to anticipate its spectral signature for incorporation into the detection algorithm. The threat cloud must also provide enough radiative signal to be detectable above the noise. Radiative transfer models predict the magnitude and spectral characteristics of the radiative signal that a plume will produce [10].

Molecules in a vapor absorb radiation in the same spectral bands at which they emit, and so the dominant radiative transfer mechanisms are absorption and emission. Some chemical agent dissemination methodologies, however, produce particles with sizes on the same scale or larger than the wavelength of the infrared light being used to detect them; in that case, the dominant radiation mechanism is scattering. Figure 2 illustrates the various radiative pathways that contribute to overall signal. The mode in which the sensor operates (active/passive) and the physical state of the threat (vapor/aerosol) determine which radiative transfer terms are dominant and which may be neglected. For passive vapor detection, for example, the term due to radiation scattered by the plume can be largely neglected. Table 1 summarizes which terms are important for each sensing mode.

Spectral Signatures

The ability to sense a chemical plume relies on a difference between radiation detected when the plume is present and not present. Thus it is the plume's scattering and absorption/emission that divulge its presence.

Detecting a chemical plume requires isolation of its spectral signature from that of the background. When we isolate the effect of the plume, the relevant measurable for passive and line-of-sight systems is not the raw measured radiance signal, but rather the differential radiance between the two conditions: plume present and plume not present. The differential signal for the vapor case depends linearly on the product of the absorption coefficient of the chemical, its concentration, the path length of the infrared radiation through the plume, and the temperature difference between plume and background, if the absorption and temperature difference are small. This linear relationship allows the use of simple matched-filter processing algorithms that compare the signal to the spectral signatures of various threat agents, providing a metric that can be compared with a predetermined threshold [11]. As with any sensing system, of course, there is a tradeoff between the probability that a threat will be detected and the probability of a false alarm—that is, setting the threshold to achieve a higher probability of detection will result in a higher false-alarm rate.

Accurate estimates of the threat-free background are a key component of the detection algorithm. For the active line-of-sight mode of operation, the plume-free spectrum is simply the hot-source spectral radiance attenuated by the atmosphere. Passive backgrounds, however, are more complicated. There are different methods of estimating the passive plume-free spectrum. For a sensor staring at an unchanging background, one may use a measurement at an earlier time, assuming that there is confidence that no plume was present at that time and temporal fluctuations of the spectrum are small. The situation becomes more complicated when the background is rapidly changing, such as when the sensor is mounted on a moving platform. When the emissivity and temperature of the background scan are not the same as for the detection scan, the signal will contain spurious background elements that may distort the spectral signature and fool the detection algorithm. To mitigate this problem, algorithm designers have developed methods to better estimate the background, and detection algorithms are used that de-weight the spectral channels

Table 1: Dominant Radiative Terms

SENSING MODE/ THREAT TYPE	SCATTERED COLD SKY	SCATTERED SOURCE	PLUME EMISSION	PLUME TRANSMISSION
Passive/Aerosol	✓	✓	✓	✓
Active/Aerosol		✓		✓
Active/Vapor				✓
Passive/Vapor			✓	✓

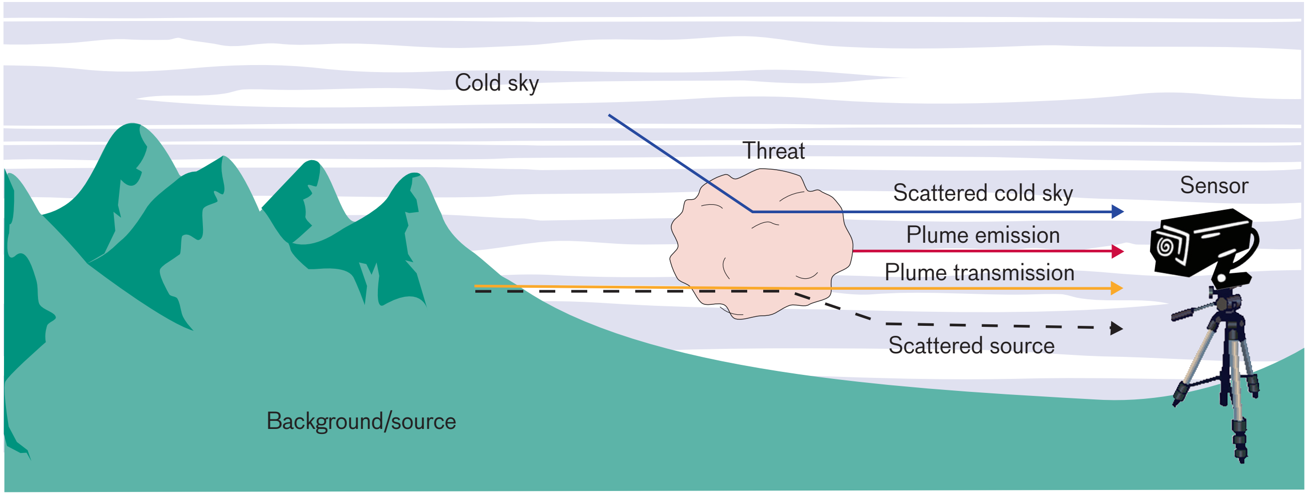

FIGURE 2. Radiance measured at the sensor is a sum of cold sky radiance scattered by the plume, plume emission, and background/source emission transmitted and scattered by the threat.

that are most affected by fluctuating backgrounds.

When we are operating a hyperspectral imager, one method of estimating the background involves averaging the spectra over all of the scene pixels. However, a scene may contain several spectrally distinct spatial subregions. An improved background estimation involves dividing the scene into spectrally similar regions and using only the most relevant subscene spectrum to estimate the background. Covariance-based detection algorithms determine which linear combinations of spectral channels contain the greatest fluctuations. Then the algorithms give less weight to readings from these channels. While this approach increases the post-processing signal-to-noise ratio, there is a limit to how much improvement post-processing can provide; detection of chemical agents in some backgrounds is inherently difficult.

Ambient vapors and aerosols can also confuse the detection algorithm. Although the atmosphere is relatively transparent in the LWIR region, atmospheric gases, which are always present, do have some features in this window. Ozone gas, for example, has an absorption feature near that of a typical nerve agent. Interferents—that is, materials that interfere with the detection of a chemical agent—can mask a real threat, or appear spectrally similar enough to a threat to cause a false alarm. In addition to disturbing the spectral observations, atmospheric absorption also attenuates the signal. Other transient vapors and aerosols, such as diesel exhaust, may interfere with threat detection as well. Because each chemical theoretically has a distinct LWIR signature, designing a sensor with high enough spectral resolution helps to differentiate between threats and interferents (although higher spectral resolution comes at the cost of lower signal-to-noise ratio).

Sensor Performance

A chemical sensor's performance parameters must be traded so as to maximize the probability that the sensor will detect a real threat while keeping false alarms at an acceptable level, and also conform to the logistical constraints of the mission. Logistical constraints include limits on volume, mass, and electrical power. For tactical applications, setup time is also important. There are several important factors that determine a sensor's ability to meet detection and area coverage requirements: noise-equivalent differential radiance (NEdN), field of view, spectral resolution, and spatial resolution. The various sensor components can be mixed and matched to achieve the optimum sensor design for the application.

NEdN indicates the detector's fundamental noise floor, and is defined as the radiance signal that produces a unity signal-to-noise ratio—the higher the NEdN, the worse the sensor's performance. If a chemical plume provides substantial radiant signal contrast from the background scene, or if the mission permits sufficient signal averaging, a high-NEdN detector may be acceptable. A situation in which substantial signal is not available or signal averaging is not feasible will require a lower NEdN (higher performance) detector. A variety of IR detector technologies are viable for chemical remote sensing,

Table 2: Selecting a Spectrometer

TYPE	SIZE	COST	SENSITIVITY TO VIBRATION	SPECTRAL RESOLUTION
Grating	Largest	Lowest	Less sensitive	Moderate
Fabry-Perot	Moderate	Highest	Sensitive	Moderate
Fourier-transform IR (FTIR)	Moderate	Highest	Sensitive	Highest
Wedge interferometric spectrometer (WIS)	Smallest	Moderate	Less sensitive	Lowest (but adequate)

including the microbolometer, pyroelectric, and mercury-cadmium-telluride (HgCdTe) devices. Microbolometers and pyroelectric detectors generally have higher NEdNs than the other options. On the other hand, they are less expensive, and because they can be operated uncooled, they consume less power. ACRs achieved with these higher NEdN devices are lower because of the need for increased signal averaging (longer dwell). HgCdTe detectors have smaller NEdNs and thus higher ACRs. However, they are more expensive and require cryogenic cooling, which increases overall sensor mass, volume, and electrical power consumption.

Instrument field of view—the area that the sensor views at any given time—is determined by the optical system, the number of FPA pixels, and the size of the detectors. The ACR depends on the field of view and velocity at which it scans the scene. A low ACR requirement permits the use of a high-NEdN detector array and signal averaging, while a high ACR requirement necessitates using either a low-noise FPA together with fast scanning or else a sensor with large field of view. To obtain favorable detection statistics, the spatial resolution should permit chemical threat clouds to fill several pixels completely. Spatial resolution can be enhanced for a fixed field of view by using an array with more pixels. Conversely, introducing more pixels at constant spatial resolution will increase the field of view.

The best spectrometer choice for a particular mission depends on multiple factors, including the required spectral resolution (Table 2). Sensor spectral-resolution requirements depend on spectral features of threat agents such as absorption linewidth. Bulky, laboratory-grade instruments can easily achieve resolution on the order of a few nanometers. However, experiments and analysis suggest that a resolution of 50 to 100 nm is sufficient to identify most known chemical warfare agents. This relatively relaxed spectral-resolution requirement allows consideration of smaller, less complex spectrometers than those traditionally used in high-resolution laboratory sensors.

Spectrometers available in the long-wavelength infrared include diffraction gratings, Fabry-Perot etalons, Fourier-transform interferometers, and interference filters. Table 2 summarizes some of the important spectrometer attributes to consider when choosing the appropriate spectrometer for a particular mission. The relative sizes reported are inclusive of the optical package paired with each spectrometer. The relative costs are based on commercially available products that are ready for integration into a sensor or could be easily adapted. However, these rankings do not necessarily capture the potential for reducing cost as production quantities grow and designs are modified.

Additionally, the spectral resolution reported in Table 2 is not a statement of the theoretical limit of these spectrometers. Rather, it gives the resolution performance in the context of practical constraints. For example, a Fabry-Perot spectrometer can theoretically achieve very high spectral resolution, but is practically limited by the imperfections in etalon surfaces. In practice, Fabry-Perot spectrometers have a resolution on the order of tens of nanometers. Sensitivity to vibration is particularly important for sensors on moving platforms. Fourier transform infrared (FTIR) and Fabry-Perot spectrometers rely on highly sensitive moving optical components and thus require careful designs to mitigate vibration effects, resulting in increasing cost and complexity. In comparison, grating or wedge interferometric spectrometers (WIS) are inherently less sensitive to vibration.

Wide-Area Chemical Sensor

One of Lincoln Laboratory's standoff chemical monitoring systems addresses the need for a compact and robust chemical sensor that can survey large areas in a short time for use on various stationary and mobile platforms. This wide-area chemical sensor, or WACS, can be used for reconnaissance missions to ensure that troops can safely enter an area, or used for site, zone, or route protection, among other missions. WACS fills an important void in chemical-sensor performance: no currently fielded sensor is optimized for both on-the-move operation and high area coverage with high spatial resolution.

Like JSLSCAD (the current military standard sensor), WACS is a passive LWIR system. But while JSLSCAD uses a single-pixel detector, WACS employs a 2-D FPA, a modification that provides high area coverage while maintaining high spatial resolution. This improvement to the current standard is owed to the increasing availability of focal planes that operate in this spectral region. WACS has the potential to be more compact and rugged than other passive hyperspectral imagers [7, 8]. Its compact nature is enabled by the use of a WIS, which has a high angular acceptance and thus can be paired with very compact optics. Most spectrometers deliver the desired spectral resolution only if incoming light is nearly collimated—a requirement that entails the use of bulky reimaging optics. Because the WIS uses a high-refractive-index substrate, its collimating requirements are much less strict than those of other spectrometers; no reimaging optics are needed, significantly reducing the size of the package. Since the WIS has no moving parts, it is easier to ruggedize for field use and requires less stabilization.

The WIS consists of alternating layers of two different dielectric films deposited on a planar substrate; the thicknesses of the layers vary along one axis, so that each spatial position along that axis transmits a different wavelength of light. When the WIS is mounted atop a 2-D FPA and paired with appropriate imaging optics, the scene is imaged; each row of the FPA views a different color of LWIR light at one particular time (Figure 3). To acquire spectral information for every spatial point in the scene, we must scan the scene along the wedge (spectral) axis. The result is a 3-D data cube with mixed spectral and spatial dimensions. If the scan is linear, a simple transformation returns the three independent spectral and spatial axes. If the scan pattern is complex, such as from a moving platform, reconstruction is more difficult and certain data may be missing. Various scanners can be used, including a folding mirror, a gimbal [12], and a Risley prism scanner [13, 14]. We have developed concept designs for several variations of this WIS-based configuration [15]. They address potential search or staring applications, and represent different trade-offs between performance, cost, size, and complexity.

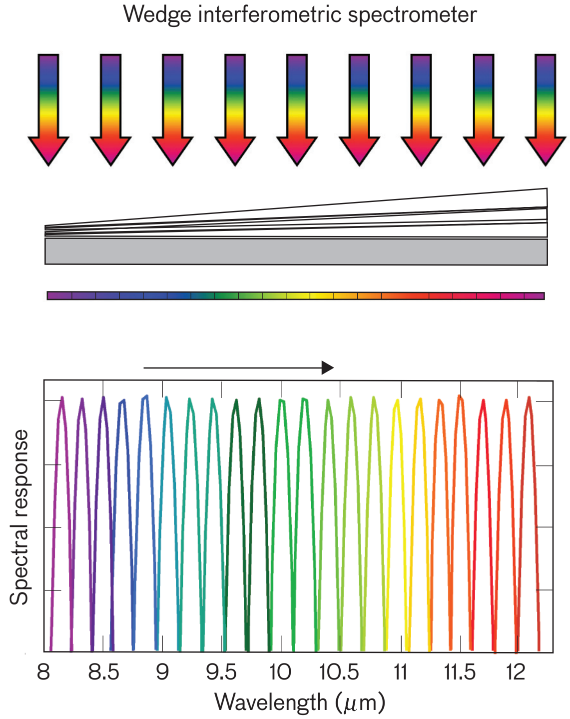

FIGURE 3. The WIS transmits a different color of infrared light at each point along its axis of linearly varying thickness. Each row of pixels of the focal plane below views the color of light transmitted by the slice of WIS above. A linear scan of the scene across this axis ensures that all colors of light are recorded for each spatial point in the scene.

Sensor Fabrication

We have constructed a laboratory testbed to demonstrate the key components of the WACS. This effort included fabrication of a WIS and subsequent mounting to a cooled, 128 × 128 pixel HgCdTe FPA in a liquid-nitrogen

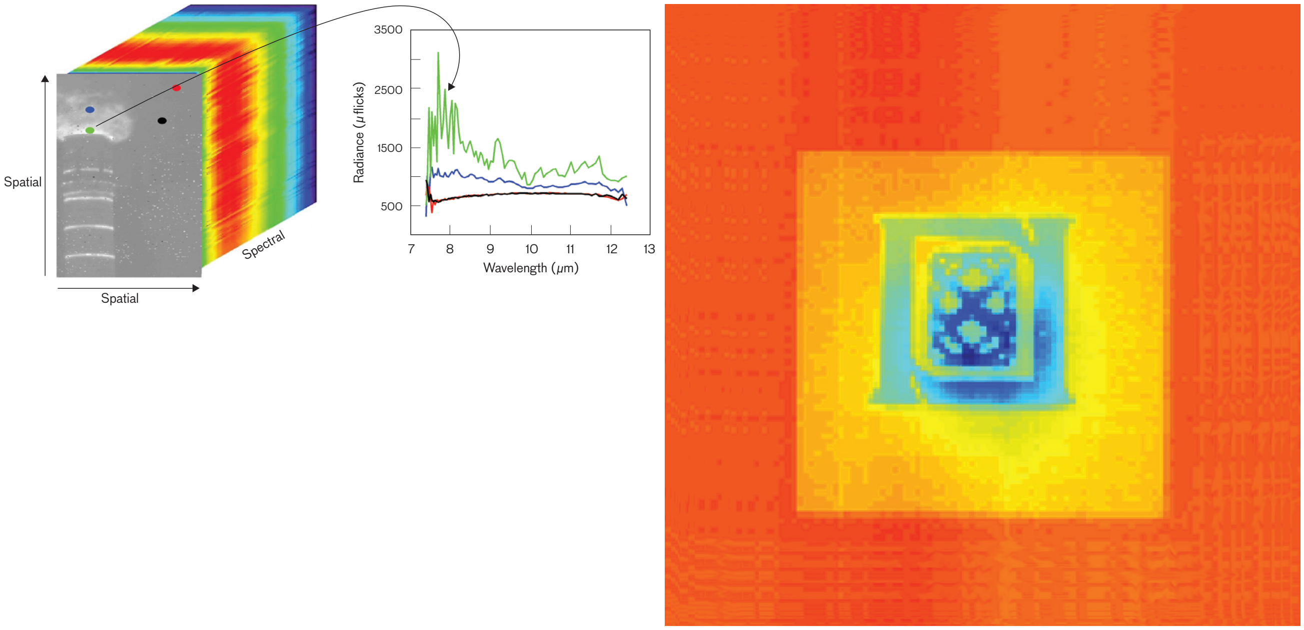

FIGURE 4. The target for proof-of-principle WACS experiments consists of a thin film of polyethylene mounted on an aluminum frame, with a superimposed Lincoln Lab logo composed of opaque black polypropylene (visible-light image, left). The false-color infrared image (right) shows a projection of the corresponding data cube onto the spatial plane.

dewar, as well as design of and integration with an optical system and scanning mechanism. We characterized the WIS/FPA pair spectrally with an Oriel MS257 diffraction-grating monochromator, and then generated sample hyperspectral data cubes. A scanning mirror provided requisite image motion for the initial experiments.

The first Lincoln Laboratory–fabricated WIS achieved an effective resolution of about 130 nm. Subsequent design and fabrication refinements have yielded an optimized WIS with 100 nm spectral resolution. The experiments presented here used the original, nonoptimal WIS test piece.

Experimental Data

To demonstrate the testbed's functionality, we acquired 3-D hyperspectral data cubes. The target scene was selected to have distinctive spatial and spectral properties that would illustrate both the imaging and spectrometric capabilities of our system. This scene consisted of a hot blackbody source (120°C) behind a sheet of low-density polyethylene (LDPE) with an overlaid mask of the Lincoln Laboratory logo. We chose polyethylene because of that material's distinctive spectral features. Figure 4 depicts a visible-band photograph of this target scene without the blackbody source, as well as a false-color infrared image of the same target acquired by WACS.

Figure 5 shows the spectrometric performance of our testbed system. The top spectrum is the low-density-polyethylene film signal averaged over a single row of pixels. For comparison, the "truth" trace at the bottom shows the equivalent spectrum acquired by a FTIR instrument that provides both higher resolution and lower NEdN [8]. The spectrum of the Lincoln Laboratory logo, also signal-averaged over a row of appropriate pixels, correctly lacks strong spectral structure due to its flat-black composition. Spectra acquired with our testbed were noisier than desired because of known damage of the FPA and use of a suboptimal WIS, but they effectively validate the baseline sensor design concept.

Inexpensive Chemical Agent Detection System

A variety of military and civilian fixed sites are at risk of chemical attack. The requirements of protecting a fixed site differ considerably from an on-the-move scenario, since there is presumably adequate time to optimize and install the detection system. Also, because there are numerous sites to protect, system cost becomes a paramount concern.

The fixed-site protection mission imposes unique requirements on sensor performance. First, wide deploy-

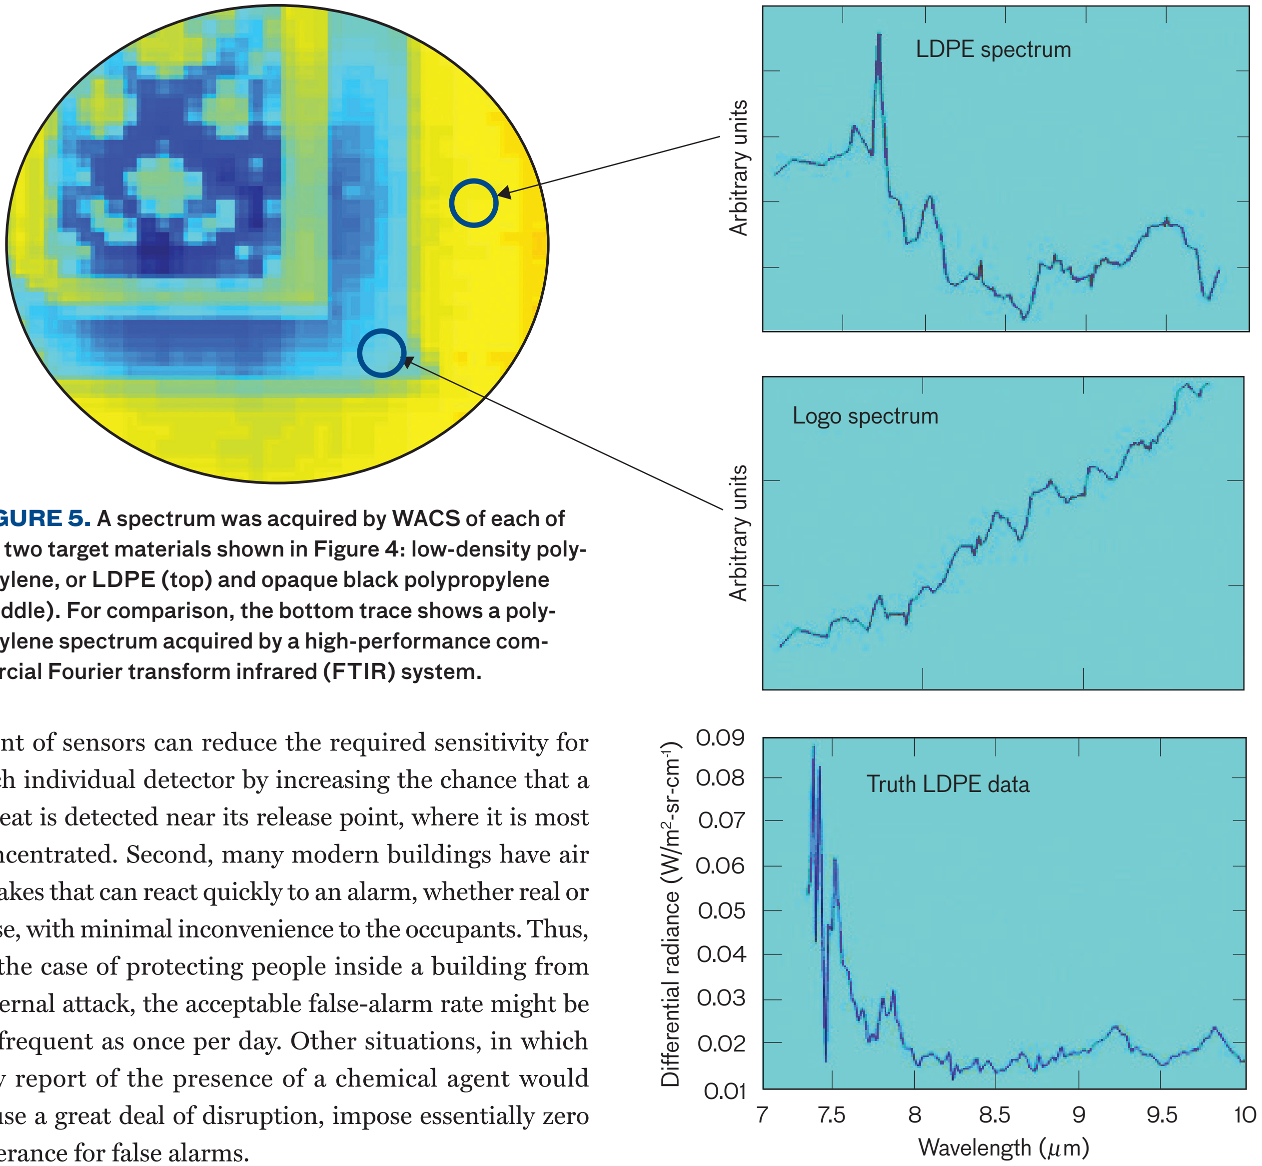

FIGURE 5. A spectrum was acquired by WACS of each of the two target materials shown in Figure 4: low-density polyethylene, or LDPE (top) and opaque black polypropylene (middle). For comparison, the bottom trace shows a polyethylene spectrum acquired by a high-performance commercial Fourier transform infrared (FTIR) system.

ment of sensors can reduce the required sensitivity for each individual detector by increasing the chance that a threat is detected near its release point, where it is most concentrated. Second, many modern buildings have air intakes that can react quickly to an alarm, whether real or false, with minimal inconvenience to the occupants. Thus, in the case of protecting people inside a building from external attack, the acceptable false-alarm rate might be as frequent as once per day. Other situations, in which any report of the presence of a chemical agent would cause a great deal of disruption, impose essentially zero tolerance for false alarms.

One way to provide fixed-site protection might be to use WACS or another passive hyperspectral sensor to continuously scan the site from a central location. However, this option is expensive, suffers from line-of-sight blockages by buildings and trees, and is limited in sensitivity by variations in the scene spectral radiance. Another method being developed by a number of groups uses an array of point chemical sensors. The difficulty here is that detecting a release with adequate warning time will require a large number of point sensors, driving up the system cost.

Lincoln Laboratory has taken a third approach, which relies on distributed, active line-of-sight detectors that can be made inexpensively while still offering adequate sensitivity to most threats over a wide geographical area. We call our detection concept the Inexpensive Chemical Agent Detection System (ICADS). ICADS consists of a number of active line-of-sight detection legs (source/detector pairs) distributed strategically around and within a site. The overall system hardware cost depends on the cost of each individual leg as well as on the number of required legs. The cost of each leg relates to its achievable range (higher-cost detectors are more sensitive, allowing for longer range).

The geometry of an ICADS system must be tailored to the site it is monitoring. Some locations, such as military bases, have well-defined and defended perimeters and may be concerned primarily with external covert attacks. Such sites could be well protected by using a fence-line geometry. Other sites may have one or more centrally located high-value buildings that would be best protected with a hub-and-spoke arrangement in which a single central spectrometer either scans or concurrently senses each leg. Sites with a large number of obstructions (e.g., buildings or trees) may require numerous shorter sensing legs. Airfields may have one longer dimension and could be best protected by a single, high-sensitivity, higher-cost sensor-source pair rather than several lower-sensitivity, less expensive units.

There are several options for implementing an individual leg of an ICADS system. One setup has the detector and source at opposite ends; however, it is also possible to have the detector and source at the same location, with the infrared signal returned by a retroreflector placed some distance away. (A retroreflector is an optical element with the feature that all optical rays are reflected nearly along the same optical path by which they arrived.) The retroflector option has the advantage of putting all the powered hardware at one end of the path, which would be particularly useful for the hub-and-spoke geometry. The disadvantage is that the cost of the retroreflector increases as a function of required range, eventually exceeding the cost of a powered source.

Our simulation shows that at a range of 280 m, we can sense a concentration-path-length of 5.5 ppm-meter with 98% probability of detection and with less than one false alarm per month.

Proving ICADS in the Field

Our work thus far has focused on determining the sensitivity versus false-alarm rate for an individual ICADS sensor leg under real-world conditions. Using this information, we can model the performance of a complete ICADS system for a particular site (work that is currently under way).

We have investigated two ICADS sensor options. One is a moderately priced FTIR system based on a commercial spectrometer with custom optics; the other is an inexpensive custom grating-based system. Both spectrometers were paired with an inexpensive custom hot source consisting of a 1000°C filament and a 20 cm parabolic reflector. These spectrometers were used in a number of different configurations to measure various system parameters.

The FTIR-based system used a Nicolet 4700 spectrometer that offers a maximum spectral resolution of 4 nm, far exceeding our requirement of 50 to 100 nm resolution. The Nicolet features a continuous dynamic alignment system that enables it to operate for long periods—months in our tests—without recalibration. The spectrometer was placed within a custom enclosure that included a 20 cm primary reflector telescope, a spectrometer, a detector, and a laptop computer. In the tests described here, the system was operated with a cooled detector element offering the highest possible sensitivity. In later measurements, the system operated with an uncooled detector element. In all tests, the system survived hot and cold ambient temperatures without failure. The overall system cost, including the cooled detector, is about $40,000. Other fully integrated and ruggedized commercial open-path FTIR systems cost in excess of $100,000. We used the Nicolet to measure the backgrounds (line-of-sight attenuation due to natural and man-made materials) in various locations at various times of year. Because of its high resolution and repeatability, the Nicolet also served as a reference sensor during tests of the grating-based spectrometer, where both systems viewed the same source.

The grating-based spectrometer system (Figure 6) consists entirely of commercially available standard parts, including the grating, a 5 cm focusing mirror, and a 64 element pyroelectric linear array [16]. Cost for these parts was less than $5000—a figure that would be reduced considerably if purchased in quantity. The spectrometer was placed in a thermally isolated enclosure during outdoor measurements. The spectral resolution of such a system is at best the spectral range divided by the number of elements in the sensor array. One important consideration in operating the system is that the spectral resolution is affected by the degree of collimation of the light that enters the spectrometer. Therefore, for a given desired spectral resolution at a given range, there is a limit to the size of the source that can be used. We use the 20 cm source at a range of 100 m to meet the resolution goal of 100 nm.

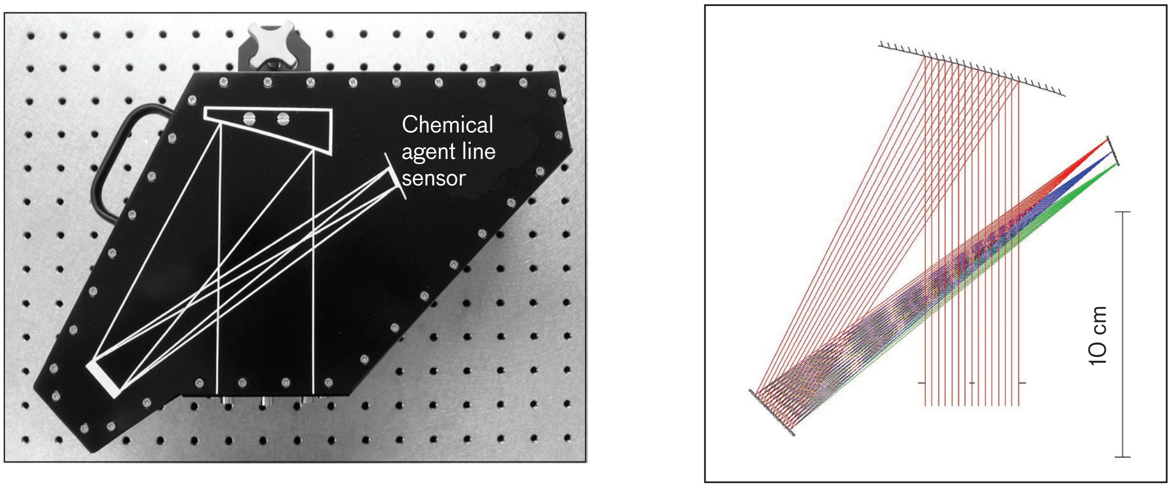

FIGURE 6. A low-cost grating spectrometer housing is shown on the left. The grating spatially separates light into its spectral components (as shown by the ray diagram on the right, where each color indicates a particular infrared wavelength); each of 64 detector elements in a pyroelectric detector array detects a different spectral band.

Because the pyroelectric detectors used in the ICADS system exhibit a significant low-frequency drift, they are typically operated in a heterodyne mode—a configuration that requires the input signal to be modulated at a frequency that is synchronized with the detection electronics. The simplest design modulates the signal with an optical chopper located within the spectrometer, where the status of the chopper can be easily communicated to the detection electronics. An alternative method that would have greatly improved system stability would have been to modulate the radiation at the source. However, this approach would have required additional engineering that time constraints did not permit. Instead, we simply accepted the constraint of using a detector that was sensitive to many different sources of radiation, including various sources within the spectrometer enclosure. Modifications since these experiments were conducted have addressed this deficiency.

Background Characterization

Overall sensor performance depends as much on the operational environment as it does on the sensor parameters. The presence of transient and ambient vapors and aerosols can resemble a threat and trigger a false alarm or reduce the overall sensitivity by attenuating the signal. (Precipitation and fog are of particular importance.) To fully understand this type of system, we made a series of measurements to assess the limitations imposed by fluctuation of ambient vapors, aerosols, and scene radiance. We took the Nicolet spectrometer to several locations—including a roadway, an indoor hallway, and an airfield—and recorded spectra for days at a time.

From these measurements we were also able to conclude that, for the 10 agents considered, the inexpensive sensor would not be affected by transient vapors, aerosols, or scene variations, but would be noise-limited. We did observe reduced sensitivity during heavy rain; the risk during rain is reduced, however, because the rain washes the chemical agent out before it reaches the intended victims. We used the ambient spectra to calculate the sensitivity versus false-alarm rate for 10 relevant threat agents. Because the attenuation spectra of most agents are known, this effect can be easily simulated in software. This simulation indicates that at a range of 280 m we can sense a concentration-path-length product of 5.5 parts per million-meter (ppm-m)—significantly lower than the lethal dose for half of the population—of a relevant threat agent with 98% probability of detection and less than one false alarm per month.

System Integration and Testing

After verifying the performance of our sensor we demonstrated both spectrometers as components of the Hanscom Lincoln Testbed (HaLT), which is under development at Lincoln Laboratory to test technologies for protecting fixed sites against biological, chemical, and radiological threats. We aimed to identify the obstacles involved in implementing such a system, to acquire long-term data in an operationally relevant environment, to develop a real-time detection algorithm, to develop the infrastructure for communication to the central command node, and to develop appropriate data fusion algorithms.

The first operational portion of our testbed was the Nicolet spectrometer and compensating source. The Nicolet system was placed on the second floor of the Lincoln Laboratory parking garage. From there it viewed the infrared source, located across the street, 280 m away, through an unobstructed path. Spectra collected by the Nicolet were transmitted via fiber-optic cable to the command post in the South Lab building, where a detection algorithm running in real time analyzed each spectrum for the presence of chemical agent. Each spectrum was compared to an average spectrum acquired prior to the current scan. The difference was then matched against a library of 10 possible chemicals, including several chemical warfare agents, toxic industrial chemicals, simulants, water vapor, and ozone. The water vapor and ozone were added to test against known atmospheric conditions. If an alarm was declared, a file was passed to the HaLT decision software, where it could be further analyzed.

The performance of the Nicolet system was impressive. During three months of operation we observed no false alarms that involved mistaken reports of the presence of any of the chemical warfare agents or toxic industrial chemicals on our list (in a demonstration of system health, we did observe several expected alarms resulting from water vapor). We again determined the system sensitivity to real agents by simulating the absorption in software, as described above. We conclude that the false-alarm rate and sensitivity of the Nicolet detection system (or one like it) will most likely meet or exceed the needs of many ICADS applications.

We have far lower expectations for the performance of the grating spectrometer than for the Nicolet, primarily because of the grating system's use of an uncooled detector array. The method of using calculated attenuations to simulate the sensor response to live agents also proved to be problematic. For the tests that were performed over a period of several weeks, the grating spectrometer was installed next to the Nicolet, where it was active during several system tests, as described below.

To evaluate our systems with a real absorbing gas, we released small quantities of sulfur hexafluoride (SF_6), a commonly used innocuous vapor that is a strong IR absorber. Release points are shown in Figure 7. During the releases the wind was blowing primarily from the south-southwest. But because the release site is in a pocket of turbulent air between a small hill and the building, instantaneous local winds fluctuated in both direction and speed. In both detectors these releases exhibited a strong, unmistakable signal much larger than the noise (the sole exception was when we released the gas downwind of the sensor line of sight).

Our next test allowed us to compare the grating response directly to that of the Nicolet. We aimed both sensors at a single source at a range of 100 m. With confidence in the Nicolet calibration, we used it as our truth sensor and as a way to calibrate for concentration path length. To ensure that both sensors viewed the same concentration path length, we released the SF_6 adjacent to the source where the detection paths converged. The response of the grating spectrometer is shown in Figure 8. When using a sensitivity setting that would produce less than one false detection per day, the grating spectrometer responded to five of the six releases. We conclude that in most cases the existing detector can detect a 7 ppm-m SF_6 challenge. We reviewed the data in the one instance in which the system failed to detect 7 ppm-m; it appears that immediately prior to the release, the detector was experiencing a large drift. This drift indicates a hardware problem that has since been addressed.

Accounting for differences in IR absorption between our simulant and a real threat, the grating spectrometer would be sensitive to approximately 33 ppm-m of a threat agent at a range of 100 m with one false alarm per day. Although this performance is, as expected, significantly worse than that obtained for the Nicolet, it still falls within reasonable limits. The grating spectrometer is a viable detection technology that has already met our requirement of costing less than $5000. We believe that with future modifications it will also meet our detection requirements.

Next Steps in Standoff Chemical Monitoring

Both WACS and ICADS are works in progress at Lincoln Laboratory. WACS development has to date focused on one particular implementation of a chemical imager,

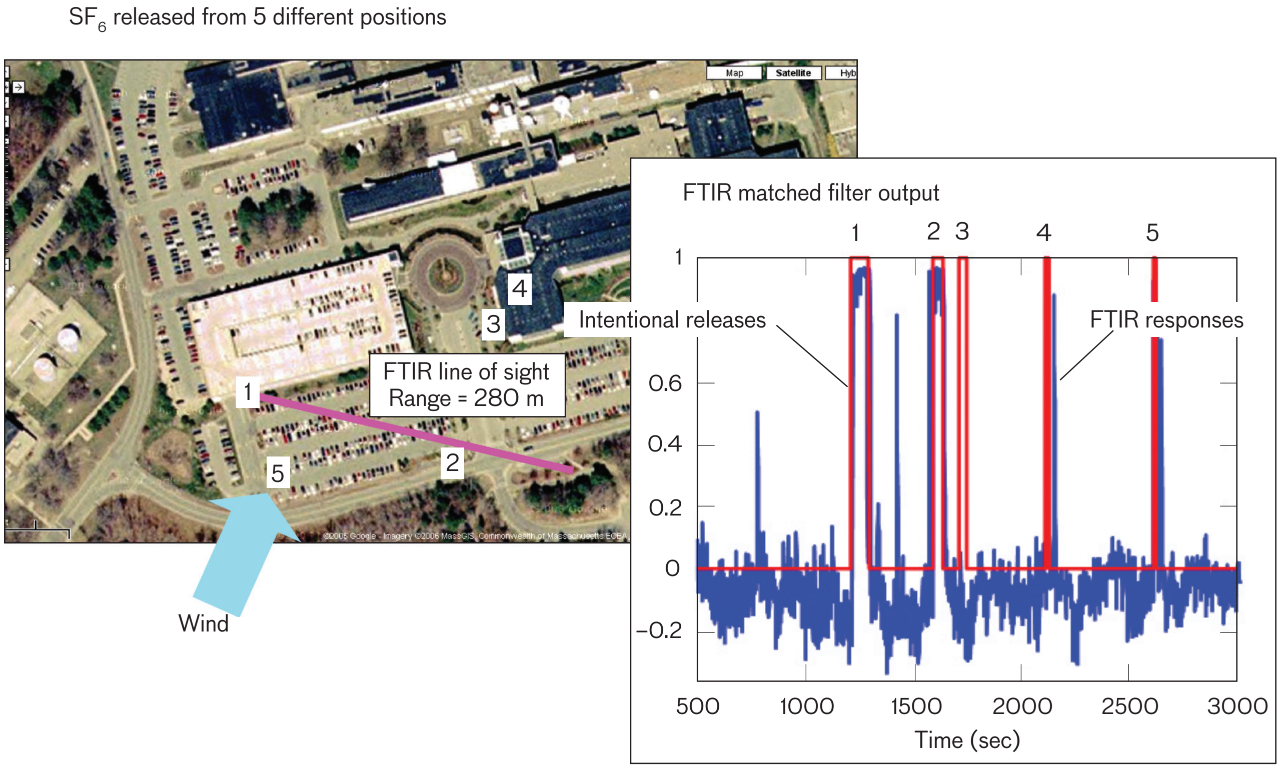

FIGURE 7. We released sulfur hexafluoride (SF_6) from five different locations at our site, as illustrated on the left. The wind was blowing from the south-southwest. The plot on the right indicates time of intentional release (red) and corresponding detector output (blue). The first was released at the FTIR spectrometer and was seen immediately. The second was released directly in the line of sight and was also seen immediately. The third release was downwind from the line of sight; as expected, we did not see it. The fourth release was at the inlet of the building's HVAC system, also downwind from the line of sight. We did not expect to see this release either—but we did (after a delay). The gas had been pulled through the building's HVAC system, expelled out the exhaust, and blown toward the line of sight. The fifth release, upwind from the line of sight, was also detected. The spike prior to any release was probably a result of a small leak during experimental preparation, and we believe that the spike between the first two releases was a result of the turbulent flow inside the canyon formed by the buildings and the hill. Both of these unexpected detections exhibit temporal characteristics indicative of a real detection: gradual onset and decay.

based on a WIS spectrometer. The Laboratory is considering other missions with different requirements, which may lead to different hardware choices. The WIS is a rugged, compact design, but reconstructing spectral and spatial data acquired from chaotically moving platforms presents difficulties. For certain missions, other spectrometers may be advantageous. We are assessing the mission requirements of two specific missions—human-portable and operation from an unmanned aerial vehicle (UAV)—and designing sensors optimized for these missions. The designs seek to leverage commercial off-the-shelf elements when applicable and are optimized for low mass, power, and cost.

A human-portable sensor would have to be battery powered. Given the high power consumption of a mechanically cooled FPA, uncooled microbolometer FPAs seem to be the only realistic option for missions lasting longer than a couple of hours. The higher sensor noise of these FPAs can be compensated by signal averaging, which is viable in a staring application. Deployment on a UAV presents a difficult challenge. Our preliminary analysis suggests that uncooled FPAs may have insufficient sensitivity for UAVs use because platform velocity is too rapid to allow for signal averaging. For UAVs that provide prime power, a mechanically cooled design using an HgCdTe FPA would be best. If platform-supplied power is unavailable, the need for batteries would impose a mass and volume penalty.

ICADS has focused on demonstration of two line-of-sight sensors and their characterization. Our future work will focus on two parallel tasks: optimization of low-cost sensor hardware and an architecture study

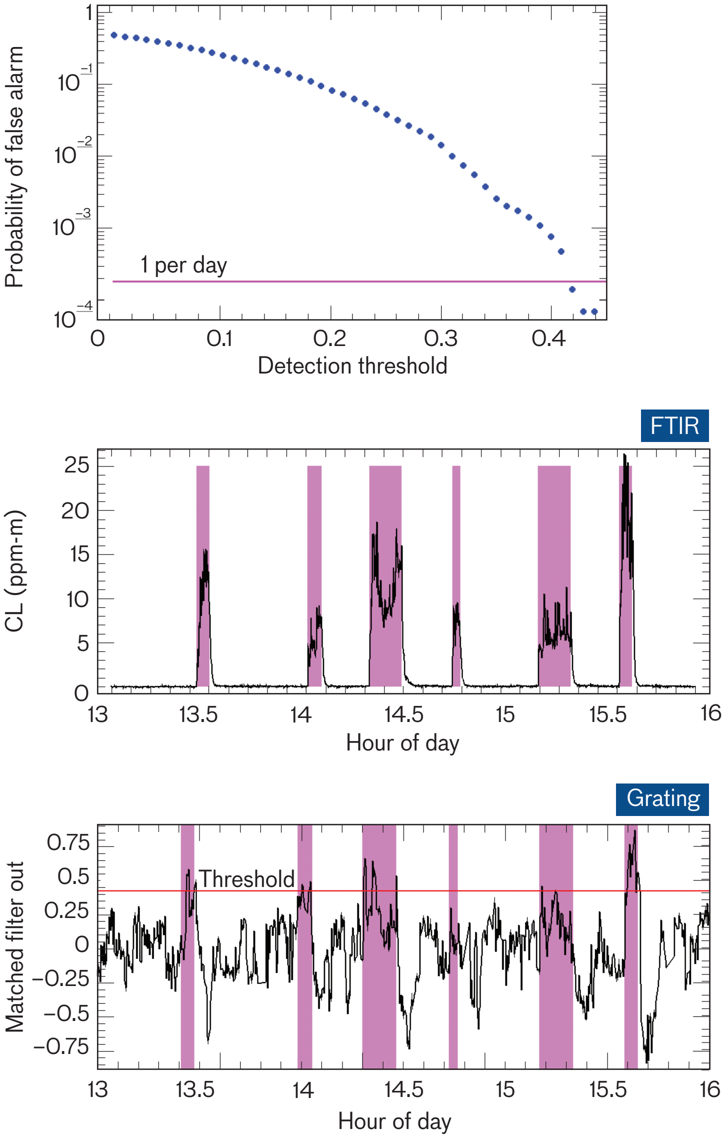

FIGURE 8. To provide truth data for the low-cost grating spectrometer, we colocated it with a high-fidelity FTIR. We released SF_6 at the source to ensure that both sensors viewed the same threat simulant. In the data from each instrument (right), the pink regions indicate when the gas was released. The concentration path length product (CL) in the middle plot is calculated from the FTIR data; the bottom plot shows the matched-filter output of the grating spectrometer. We saw a response for all six releases on the grating spectrometer, though one was below threshold. As expected, the noise was much higher in the grating system. For each release, we observed a positive spike corresponding to the threat detected, followed by a negative spike indicating that the threat had cleared. More sophisticated data processing and temperature stabilization should provide better performance. The plot at the top shows that the probability of false alarm declines as the detection threshold for triggering an alert is increased.

to quantitatively address the best placement of sensors. Sensor optimization entails deciding upon the most efficient investment of effort, either improving the existing grating spectrometer developed by Lincoln Laboratory and transferring that technology to industry or working with industry to develop a sensor to meet our cost goal. A further potential expansion of this work involves the use of quantum cascade lasers, rather than hot filaments, as the active infrared source. Such lasers can provide a higher signal-to-noise ratio, which results in a longer range, smaller optics, or better detection performance.

The sensor placement study will begin with an examination of several canonical scenarios, including different sites, with different levels of obstruction, different threats, different sensor types, and different geometrical configurations. We will leverage the data acquired by the Nicolet and grating spectrometer systems to create realistic sensor models. These data will allow us to observe scaling laws and physical trends with regard to sensor placement, the number of required sensors, and sensor range. The insights thus gained will help us to develop a method for designing a site-specific architecture. This process will be applied to a real and operationally relevant site—Hanscom Air Force Base, where Lincoln Laboratory resides. This optimized protection system will be implemented at Hanscom and integrated with the larger Hanscom Lincoln Testbed as an early warning chemical sensor.

Acknowledgments

Many people at Lincoln Laboratory contributed to this work. The WACS concept was conceived and initiated by Edward Wack (now director of future acquisition at the Joint Program Executive Office for Chemical and Biological Defense) and Michael Kelly. Peter O'Brien designed and fabricated the WIS. Danette Ryan-Howard and Darryl Weidler designed the opto-mechanics. Other major contributors to WACS include Curtis Colonero, John Lakness, Joseph Costa, Mark Gregg, David Kusinsky, and Robert Ashworth. The ICADS program was conceived by Ronald Hoffeld and William Herzog and first led by Thomas Jeys. Roshan Aggarwal developed the Nicolet FTIR system. Matthew Angel developed the grating spectrometer system. Other participants include Antonio Sanchez, William DiNatale, and John Aldridge. We thank David Weitz for his assistance in preparing this manuscript and for many helpful discussions. We also thank the Director of Defense Research and Engineering for sponsoring this work. ■

REFERENCES

1. D.A. Skoog, F.J. Holler, and T.A. Nieman, *Principles of Instrumental Analysis* (Saunders College Publishing, Philadelphia, 1998).
2. F.M. D'Amico, O.K. Emge, and G.J. Roelant, "Outdoor Chamber Measurements of Biological Aerosols with a Passive FTIR Spectrometer," *Proc. SPIE,* vol 5268, 2004, pp 173–183.
3. H. Lavoie, E. Puckrin, J.M. Theriault and F. Bouffard, "Passive Standoff Detection of SF_6 at a Distance of 5.7 km by Differential Fourier Transform Infrared Radiometry," *Appl. Spectrosc.,* vol. 59, no. 10, 2005, 1189–1193.
4. J.N. Pawloski and D.G. Iverson, "Use of Optical Remote Sensing Techniques to Monitor Facility Releases," *Hydrocarbon Processing,* vol. 77, no. 9, 1998, pp. 125–130.
5. W.G. Fately, R.M. Hammaker, M.D. Tucker, et al. "Observing Industrial Atmospheric Environments by FTIR," *J. Mol. Structure,* vol. 347, 1995, pp. 153–168.
6. B. Hammond and M. Popa, "Overview of the Joint Services Lightweight Standoff Chemical Agent Detector (JSLSCAD)," *Proc. SPIE,* vol. 5795, 2005, pp. 86–95.
7. W.J. Marinelli, M.L. Finson, J. Hagge, et al., "Wide Area Detector for Integrated Early Warning," *Proc. 6th Joint Conf. on Standoff Detection for Chemical and Biological Defense, Williamsburg, Va., Oct. 25–29, 2004.*
8. V. Farley, C. Belzile, M. Chamberland, et al., "Development and Testing of a Hyperspectral Imaging Instrument for Field Spectroscopy," *Proc. SPIE,* vol. 5546, 2004, pp. 29–36.
9. Special issue on spectral imaging, *Lincoln Laboratory Journal,* vol. 14, no. 1, 2003, www.ll.mit.edu/news/journal/14_1.html.
10. A. Beil, R. Daum, G. Matz, R. Harig, and G. Matz, "Remote Sensing of Atmospheric Pollution by Passive FTIR Spectrometry," *Proc. SPIE,* vol. 3493, 1995, pp. 32–43.
11. D. Manolakis and F.M. D'Amico, "A Taxonomy of Algorithms for Chemical Vapor Detection with Hyperspectral Imaging Spectroscopy," *Proc. SPIE,* vol. 5795, 2005, pp. 125–133.
12. Sagebrush Technology, Inc., Model-30 Direct Drive Pan-and-Tilt Gimbal, www.sagebrushtech.com.
13. R.M. Marino, W.R. Davis, G.C. Rich, et al., "High-Resolution 3D Imaging Laser Radar Flight Test Experiments," *Proc. SPIE,* vol. 5791, 2005, pp. 138–151.
14. G.F. Marshall, "Risley Prism Scan Patterns," *Proc. SPIE,* vol. 3787, 1999, pp. 74–86.
15. J.A. Seeley, M. Kelly, E. Wack, et al., "An Advanced Wide Area Chemical Sensor Testbed," *Proc. SPIE,* vol. 5995, 2005, pp. 59950W-1–59950W-11.
16. J.A. Seeley, M.M. Angel, R.L. Aggarwal, et al., "Inexpensive Chemical Defense Network for a Fixed Site," *Int. Symp. on Spectral Sensing Research (ISSSR), Bar Harbor, Maine, May 29–June 2, 2006.*

ABOUT THE AUTHORS

Juliette A. Seeley is a staff member in the Sensor Technology and System Applications group. She received her doctorate in physical chemistry from the University of California at Berkeley, where she helped develop novel nuclear magnetic resonance techniques. Her work at Lincoln Laboratory focuses on standoff chemical sensing.

Jonathan M. Richardson is a staff member in the Sensor Technology and System Applications group, where he primarily works on chemical and biological countermeasures. He received his bachelor's degree from the University of Pennsylvania and his doctorate in physics from Harvard University. He performed his thesis research at the National Institute of Standards and Technology's Center for Neutron Research, where he made precision measurements of neutron decay.