Health Surveillance and Diagnosis for Mitigating a Bioterror Attack

Adam Szpiro, Bernadette Johnson, and David Buckeridge

The most significant factor in minimizing the adverse impact on public health of a biological-warfare-agent aerosol attack is the time required to receive appropriate treatment. We look at two human-health surveillance techniques designed to reduce the time window between attack and the start of treatment. In the first technique, syndromic surveillance, data are assembled from a variety of sources, including primary reported symptoms in emergency departments, calls to 911, pharmacy records of specific treatments, and school absenteeism. Assimilation of these data sources can provide an early indicator of a regional or local outbreak of infectious disease. The second approach, the Biological-Agent Correlation Tracker (BACTrack), offers a means of locating, in space and time, the probable origin of an attack through the use of a volunteer population who report their health status to a central source and who also carry some kind of location tracking device.

» **Lincoln Laboratory has been designing** and developing bio-agent detection systems for both the Department of Defense and the Department of Homeland Security. Most of our work has focused on rapid environmental detection of a bio-aerosol attack. The attack information is used in some cases to initiate protective responses such as evacuation or masking, and in others to direct appropriate treatment to the targeted population. For the most part, these systems are limited to use in specific locations, against specific threat scenarios and targets. Thus, until environmental detectors are widely deployed and integrated with the public health community, the first indication of a bio-agent attack is likely to be via exposed and infected individuals presenting at points of care.

Events following the release of *Bacillus anthracis* (the causative agent for anthrax) in the United States mail system in 2001 illustrate the importance of clinical case findings in sensing and characterizing a bioterrorism incident. Ultimately there were 11 confirmed or suspected instances of inhalational anthrax, and all were identified through clinical diagnoses [1]. Information from the initial diagnoses led to an unprecedented antibiotic prophylaxis campaign involving about 10,000 individuals—an effort that probably prevented many additional cases of disease [2].

In general, how quickly a person receives a correct diagnosis depends upon what symptoms are evident, how astute the treating physician is, and what access the treating facility has to advanced diagnostics. Anthrax's initial presentation can closely resemble that of common respiratory viruses, making the disease difficult to distinguish from flu-like illness [3]. Unless there were an unusual

clustering of cases, or some other information leading to suspicion of an attack, days could elapse before the first cases were correctly diagnosed. Once even a few cases (perhaps even one) were definitively shown to result from bio-agent exposure, however, actions would be initiated to determine the probable origin and to establish whether the exposure was of natural or man-made cause.

Given the possibility of a delay in diagnosing the initial cases and determining the source of exposure, it is likely that many infected individuals would receive treatment too late and that large numbers of people who were neither exposed nor infected would receive treatment. The public-health impact of an attack can be measured in the lives lost by lack of timely treatment, the adverse health effects associated with treating well individuals, the cost of acquisition and distribution of treatment to the population, and the economic and social disruption that such a campaign would engender.

Timeline for Clinical Diagnoses

We have performed analyses to quantify the effect of detection and identification time on the outcome of a bio-agent attack. We have selected *Bacillus anthracis* as an example agent for four reasons: it is regarded as a significant bioterrorism threat, it is amenable to treatment, data exist on human exposure, and we have historical data on treatment efficacy and distribution logistics.

Rapid diagnosis of the initial anthrax cases is critical in order to enable a timely medical response that can save the lives of individuals who are infected but not yet sick. Antibiotics are highly effective at preventing death if treatment is initiated prior to the onset of symptoms. Because anthrax has an incubation period that ranges from several days to more than a month, there is a window of opportunity after an attack during which mass distribution of antibiotics can dramatically reduce the number of fatalities. In a companion article in this issue (page 115), Diane Jamrog and colleagues present the results of analyses that predict how effective such a prophylactic campaign could be as a function of how soon it commences. Their results indicate that to decrease the number of fatalities by 95%, prophylaxis must begin less than three days after an attack, assuming an aggressive timeline to complete distribution to the entire population in the course of three days. The time available to initiate prophylaxis is even less if it takes longer to reach the entire population.

To find out whether traditional medical diagnosis provides sufficiently early warning for a mass prophylaxis campaign to be effective, we employed a mathematical model that focuses on microbiological blood culture as the primary means of diagnosis and that simulates the time-evolving state of each infected individual in terms of disease progression, health-care-seeking behavior, and diagnostic-test status. Because inhalational anthrax is rare in the United States, diagnosis of a single case without a known environmental risk factor would be sufficient to conclude with high confidence that an attack has occurred. Still, a single case would not be sufficient basis on which to justify a mass prophylaxis response. We regard 10 as a representative number of cases to declare that an attack is large scale. Depending on the size of the attack, our model indicates that the first diagnosis would occur two to three days after the attack, with the tenth diagnosis following approximately one day later. As Jamrog shows, this timeline would severely limit the response effectiveness. Thus current gold-standard diagnostic assays are too slow to be of much help in characterizing a large-scale attack.

Experimental diagnostic techniques that could provide reliable identification data within minutes to hours do exist, but they are not available to the general medical community. In a study commissioned by the Defense Advanced Research Projects Agency and the Defense Threat Reduction Agency, Lincoln Laboratory examined the feasibility of building a Department of Defense health surveillance and biodefense system (HSBS) for protection against and mitigation of biological assaults (see sidebar, facing page). One objective of the study was to assess the feasibility of developing rapid, broad-spectrum diagnostic tools for detecting and identifying pathogenic microorganisms. The 2002 HSBS summary report concluded that there were no obvious technological, logistical, or legal barriers to the development of such a system.

Indeed, because emerging diagnostic technologies such as gene chips (DNA microarrays) and protein microarrays offer a way to provide thousands of concurrent assays in a single test format, it seemed reasonable to assume that investments in cost reduction or readout simplification could lead to practical diagnostic devices with the desired characteristics. The diagnostic technology would also offer direct patient benefit, hastening its acceptance by the medical community. The creation of a

The Health Surveillance and Biodefense System

A recommended approach to defending against bio-attack has yet to be implemented.

It has been more than six years since a Defense Science Board task force released the results of a study that examined deficiencies in the nation's readiness against a biological attack [a]. The study's three top recommendations were to develop a database of signatures of bio-agents that cause human disease; create a diagnostic device (Z-chip), to be used during routine clinical sampling, which would provide immediate diagnoses of diseases in the database; and set up a warning and communication system that would alert military and civilian health-care organizations in the event of a confirmed bio-attack.

A follow-on study requested by the Office of the Secretary of Defense on the feasibility of a health surveillance and biodefense system (HSBS), led by Darryl Greenwood at Lincoln Laboratory, conceptualized such a system and proposed a feasibility demonstration. The health-care provider interacts with the patient with technologies that improve speed, accuracy, and specificity. Data at the local point-of-care level are entered into terminals with patient identity protected. Information assurance implements measures to protect privacy and to protect the system from intrusion. The network can be thought of as a secure web, accessing data from a variety of sources. Data are mined in various ways and presented to a command and control center, which has the responsibility of informing authorities such as the Centers for Disease Control or the White House. Data are fed back through the system down to the point-of-care level to aid in diagnosis. The proposed system would provide both direct patient benefit as well as generating real-time, assured disease-surveillance information.

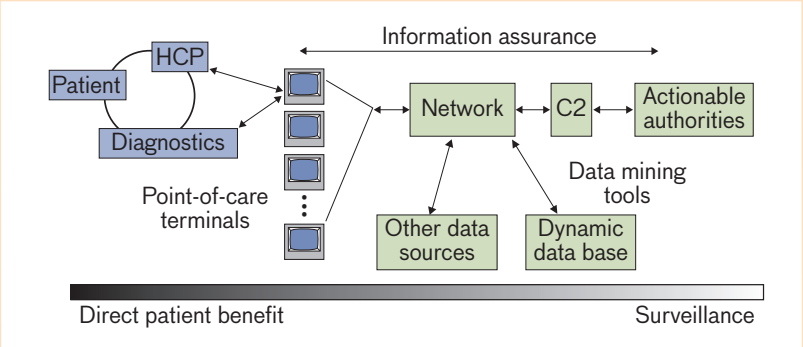

FIGURE A. In a system such as that proposed by Lincoln Laboratory in 2002, data mined in various ways are presented to a command and control (C2) center that then informs authorities such as the Centers for Disease Control. Data are fed back to the point-of-care level to aid in diagnosis by health care providers (HCP).

The HSBS is theoretically feasible. The fundamental technical barrier to its development was (and remains) the lack of rapid, broad-spectrum diagnostic assays. Today's diagnostic capabilities are largely symptom driven and rely on culturing; answers are unavailable for hours to days to even weeks. Treatment is often initiated, and even completed, without certain knowledge of the infectious organism. The few diagnostic tests now in use at the point of care (e.g., rapid immunoassays for strep) offer an excellent starting point for HSBS, but there are no technologies sufficiently mature to diagnose more than a few diseases in a multiplex fashion. In addition, the output of the few point-of-care tests that are available now is typically captured only within the health-care provider's own information network, limiting the ability to identify emerging infections across a population.

REFERENCE

a. "Protecting the Homeland: Defense against Biological Warfare: Report of the Defense Science Board 2000 Summer Study," Defense Science Board, Washington, D.C.

network that not only collected diagnostic and other disease data for attack detection but that would significantly raise disease-surveillance reporting rates and facilitate rapid response was considered an important feature in a bio-attack mitigation strategy. In short, the study endorsed HSBS as a valid, workable concept and strongly recommended that the Department of Defense undertake a program to demonstrate the technology.

Unfortunately, with the exception of an Air Force–funded pilot demonstration program called Epidemic Outbreak Surveillance [4], little has been done to advance the HSBS concept. A few infectious-disease reporting systems have been developed, but these systems rely on existing diagnostic and confirmatory assays for common (or newly emerging) diseases. They do not provide for timely two-way reporting (up to public-health authorities and back down to the point of care), and their latency is highly variable. The private sector has little incentive to develop diagnostics for biowarfare agents (as they are considered rare diseases) and the civilian sector has neither the charter nor coordination to implement a wide-scale surveillance-based information-technology network.

Two Lincoln Laboratory activities are designed to reduce the time between exposure and initiation of correct treatment. The first is an analysis of the role of syndromic surveillance, in which early reporting of common symptoms and patterns of therapeutic purchases can be used to characterize disease outbreaks. Data are assembled from a variety of sources, including primary reported symptoms in emergency departments, calls to 911, pharmacy records of specific treatments, and school absenteeism. The assimilation of these data sources can be used to provide timely information about a regional or local outbreak of infectious disease. The second approach, the Biological-Agent Correlation Tracker (BACTrack), offers a means of locating in space and time the probable origin of an attack through the use of a volunteer population who report their health status to a central source and who also carry some form of geolocation tracking.

Syndromic Surveillance

The idea behind syndromic surveillance is to monitor population patterns of nonspecific symptoms to detect evidence of a bio-agent-induced epidemic. In part because initial symptoms of many bio-agents are nonspecific and can appear flu-like (e.g., elevated temperature and respiratory distress), there is a significant delay between the onset of symptoms and the time at which cases are diagnosed. A single person exhibiting respiratory symptoms may not be immediately identified as an anthrax victim, for example, but a sudden and unexpected increase in individuals with such symptoms throughout the population could be taken as an indication of a bioterrorism attack. A syndromic-surveillance alert is based on patterns of nonspecific symptoms throughout the population and does not depend on even one individual being positively diagnosed with anthrax. The goal of syndromic surveillance is to use combinations of multiple data sources to look for such patterns. Lincoln Laboratory's primary interest is evaluating syndromic surveillance as part of an integrated system to defend against a bioterrorism attack, specifically focusing on detect-to-treat defense against an aerosol anthrax attack.

Many techniques have been devised to collect and process data that could be useful in bioterrorism defense [5]. Still, many in the public-health community remain skeptical about syndromic surveillance [6], for understandable reasons. One problem is that the information such a system provides tends to be nonspecific. Therefore, if a syndromic-surveillance system detects an anomalous pattern in population health, it would be difficult to propose initiating a mass antibiotic campaign without independent confirmation that the outbreak is caused by anthrax. However, syndromic-surveillance detection could at least cue physicians to be more alert to anthrax and other bioterrorism-related diseases. This approach holds promise for speeding detection by clinical diagnosis, but it still depends on a syndromic-surveillance alert occurring before the first clinical case finding.

Moreover, syndromic surveillance has a potentially important role in characterizing an attack after it has been detected. Diagnosis of a single case of anthrax would most likely be sufficient to conclude that there has been an attack, but it would not provide adequate information to determine that mass prophylaxis is the appropriate response. The default option in that circumstance is to delay the response until there are additional diagnosed cases. However, our analyses indicate that waiting for even the tenth case to be diagnosed could cause an additional delay of one day. Syndromic surveillance could augment medical diagnosis by providing information about the size of an attack.

We have conducted a modeling simulation study to assess the value of syndromic surveillance for early detection [7] and for attack characterization [8]. To conduct this evaluation, we require background syndromic data and simulated data corresponding to an anthrax attack. We start with authentic background data from the Department of Defenses' Tricare health-care system in the Norfolk, Virginia, area, and we tabulate the clinic visits with upper-respiratory symptoms under the ICD-9 system of classification (ICD-9 stands for International Classification of Disease, 9th Revision). We construct attack data by adding syndromic records resulting from patients in a simulated attack being assigned ICD-9 codes that correspond to upper-respiratory distress such as runny nose, cough, and sore throat. We consider a syndromic-surveillance system that searches for anomalous patterns of ICD-9 codes for such symptoms as an indicator of a biowarfare attack.

A key step in detecting or characterizing an attack is to compare the observed syndromic records to what would be expected without an attack. Our algorithm generates the expected number of syndromic visits by forecasting the historical trend by using a seasonally adjusted, autoregressive integrated moving average (SARIMA) statistical model.

Early Detection

We begin by evaluating the capability of a temporal algorithm to detect an anthrax attack and comparing the detection time to the first clinical case finding. We construct a Monte Carlo simulation of the time series of outbreak-related upper-respiratory syndromic records and add it to the authentic background data to obtain a time series corresponding to an attack. Figure 1 shows an example time series resulting from an attack affecting 50,000 people. The total number of records including anthrax cases evidently exceeds the forecast values, so we expect that a temporal-anomaly detection algorithm would generate an alert. We employ a cumulative-sum detection algorithm tuned to give about one false alarm per month, and test it on simulated outbreaks of different sizes.

For attacks affecting 10,000 people, the algorithm detects the outbreak about 60% of the time; that value increases to 100% for attacks affecting 30,000 people or more. That is a good result, but keep in mind that our primary interest is in the timeliness of syndromic-surveillance detection compared to the first clinical diagnosis. By that measure, syndromic surveillance still falls short. Even with a large attack, during which this approach has its best performance, temporal syndromic surveillance typically does not detect the attack prior to the first case finding. Note, however, that these results represent the median performance, and in some of the Monte Carlo runs, syndromic surveillance does provide an early-detection advantage.

One approach to improving syndromic surveillance is to include spatial information with the syndromic records and to search for anomalous patterns of nonspecific disease in spatial-temporal clusters. The idea is that an aerosol anthrax attack may involve releasing the agent either inside a building or from a specific area outdoors, and in either of these scenarios the exposed population would be concentrated in a small area. By looking for geographically focused disease clusters, we could rule out many of the anomalies that cause false alarms for pure temporal surveillance.

To evaluate the performance gained by adding spatial information, we extended the simulation study to include a spatial-anomaly detection algorithm using home zip codes and Poisson regression within each zip code. As Figure 2 makes clear, spatial-temporal surveillance offers an improvement over temporal surveillance, especially when the release occurs from a single point in space. In the median case, the syndromic-surveillance detection precedes the first clinical diagnosis for all size attacks considered, with a maximum advantage of a half day. Because we rely on home zip codes, these results are valid if an attack occurs at a time when most people are at home (e.g., early morning). The situation is significantly more complicated if we allow for the possibility that people are not at home when the attack occurs.

Attack Characterization

Our analysis demonstrates that syndromic surveillance has limited utility for early detection of an anthrax attack—in the median, syndromic surveillance can detect an attack one half day sooner than clinical diagnosis when the release occurs from a single point in space, and less than that for a moving line release (which might be employed by a terrorist to expose many people to a small dose while evading a bio-agent detection system). Still, the informa-

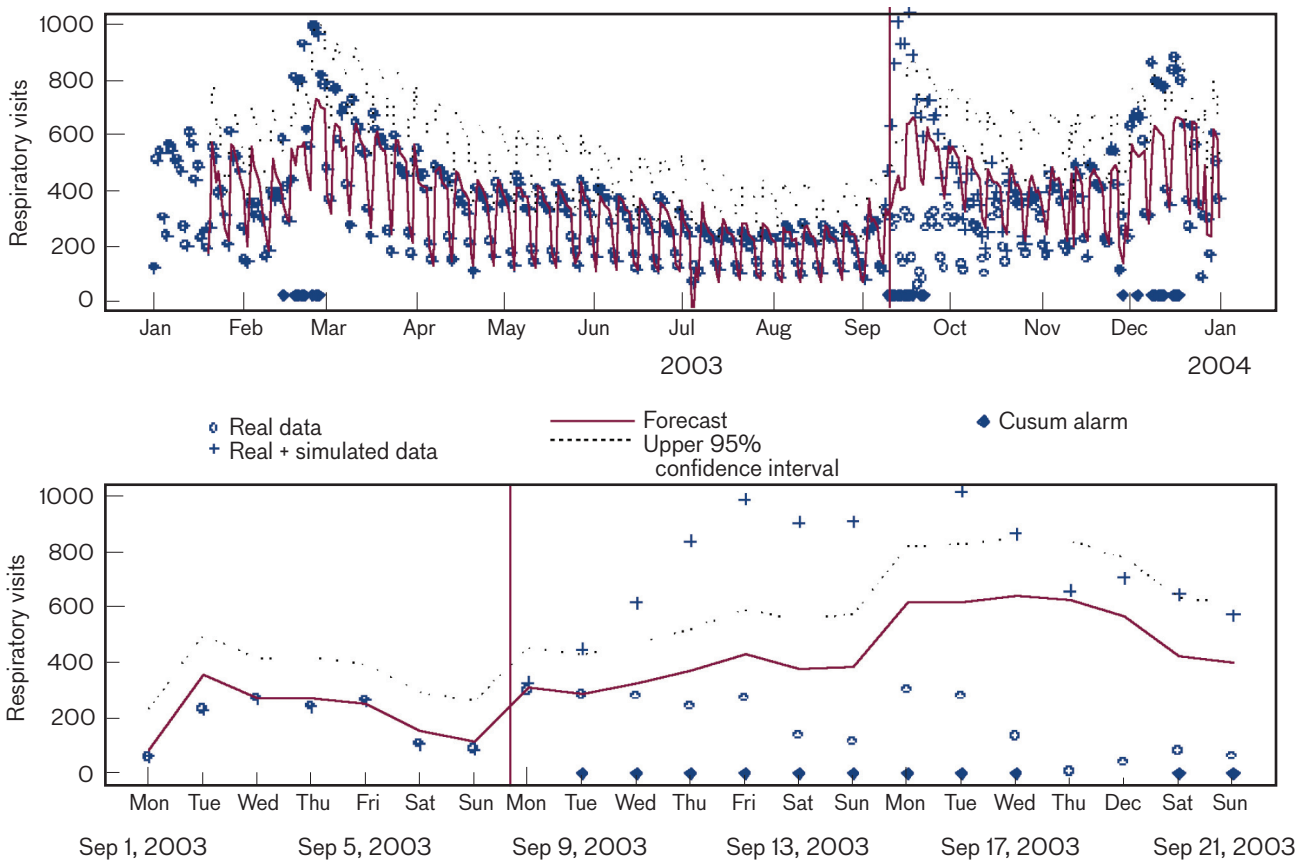

FIGURE 1. A simulated anthrax attack affecting 50,000 individuals in the Norfolk, Va., area yields data on the potential effectiveness of syndromic surveillance. The top graph shows upper-respiratory syndromic records from the simulated attack added to authentic data from the Norfolk Tricare medical system. This attack would be detected by the cumulative sum (cusum) temporal surveillance algorithm because, based on the observed trend before the attack, the total (real plus simulated) records significantly exceed the forecast. The bottom graph zooms in on days surrounding the simulated attack on September 7, 2003.

tion collected in syndromic surveillance systems could play an important role in characterizing the attack. Once it is known that an aerosol anthrax attack has occurred, the most pressing question is to determine if the attack is sufficiently large scale to warrant mass prophylaxis or if it is instead an isolated incident that should be dealt with in a more limited way. Syndromic surveillance is well suited to answering this question because it specifically tracks macroscopic population trends based on symptoms that might otherwise not be connected with the attack. The lack of specificity is a handicap for detecting an attack, but once an attack has been detected and confirmed through other means, it is reasonable to associate any statistically significant syndromic pattern with anthrax and to use that information in planning a response.

Using a simulation study similar to the one described previously for evaluating early detection, we evaluate temporal syndromic surveillance for attack characterization. Our goal is to estimate the number of patients with symptoms that can reasonably be attributed to the attack (e.g., upper-respiratory distress). This number—called the excess syndromic burden—can set a lower bound on the scale of the attack. One could contemplate extrapolating further to estimate the total size of the attack [9]. Because of uncertainties about the course of the disease and variations from strain to strain of anthrax [10], however, we prefer to focus on estimating the more fundamental quantity of excess syndromic burden.

The excess syndromic burden is estimated by subtracting the projected number of syndromic records from the number observed in the days immediately preceding detection. The syndromic data for our study are aggregated daily, and we initially assume that attack characterization needs to occur immediately upon diagnosis of the

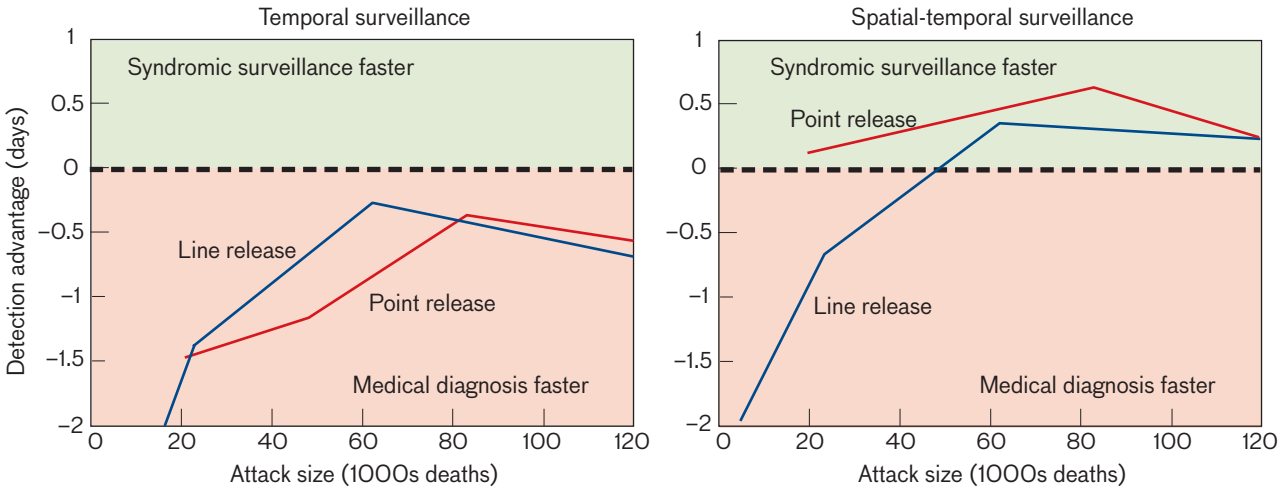

FIGURE 2. Under some circumstances, a bio-attack could be detected more quickly by using syndromic surveillance than by relying on medical diagnosis. This advantage in the median case comes, however, only for large-scale attacks and only if the syndromic surveillance involves temporal as well as spatial monitoring. For all attack sizes, temporal syndromic surveillance is less timely.

first case. Thus only data collected through the beginning of the diagnosis day are available.

Figure 3 illustrates how the difference between projected and actual records could be used to estimate the excess syndromic burden for an attack that affects 50,000 people. The first clinical case finding occurs 4.5 days after the attack, so we use syndromic data from the end of the fourth day to estimate the excess syndromic burden attributable to the attack. In this case, the actual number of syndromic cases from the attack is 245, and the estimated number is 263. We repeat this calculation for multiple Monte Carlo runs for attacks ranging from 1000 to 50,000 untreated fatalities. Figure 4 shows that the median estimate for all attack sizes is less than 100 and is not statistically significant; i.e., the observed estimates would not be atypical even in the absence of an attack, and are thus of limited value.

Something else is also evident from Figure 4: if we wait until the end of the diagnosis day and use syndromic data collected on that day, the estimates tend to be larger and are marginally statistically significant for attacks affecting 30,000 people or more. One difficulty with this approach is that additional people might seek care after hearing about the initial diagnosis, leading to a bias in the estimated attack size. Another drawback is that it entails

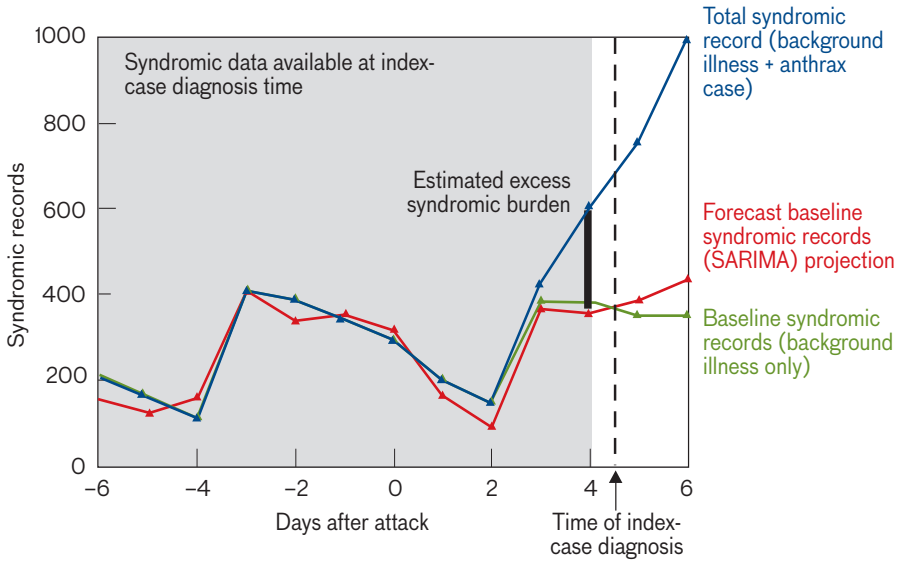

FIGURE 3. Temporal syndromic surveillance could help in characterizing an attack that has already been detected through other methods. In the case of this Monte Carlo run for an attack affecting 50,000 people, the first clinical case finding occurs 4.5 days after the attack. Syndromic data are available through the end of the fourth day and can be used to estimate the excess syndromic burden attributable to the attack by comparing the observed number of records to the forecast number. SARIMA is the seasonally adjusted autoregressive integrated moving average. The actual number of syndromic cases from the attack is 245, and the estimated number is 263.

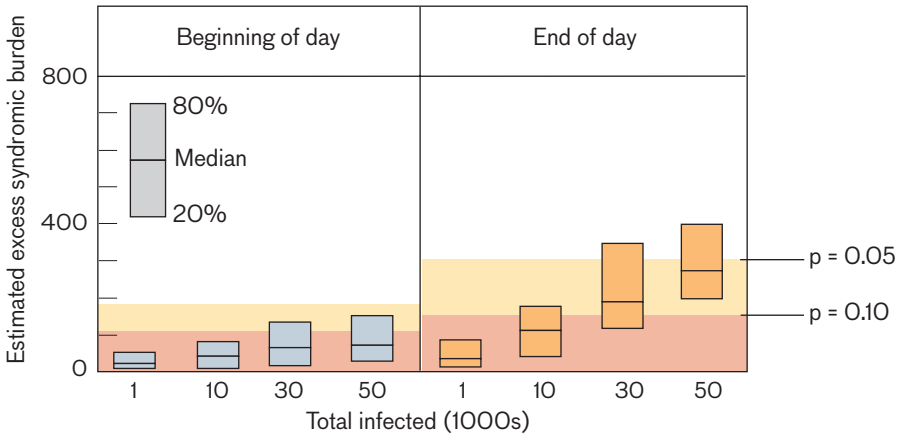

FIGURE 4. This syndromic-surveillance simulation shows the median values and 80/20 percentile for the estimated excess syndromic burden attributable to an anthrax attack. When characterization uses syndromic data accumulated through the beginning of the day on which the first clinical diagnosis occurs (left), the median values are consistently less than 100 and are not statistically significant (as defined by p 0.10). When characterization uses data accumulated through the end of the diagnosis day (right), the estimates tend to be larger and, for attacks affecting 30,000 people or more, are marginally statistically significant (0.05 < p < 0.10).

an additional delay in characterizing the attack beyond the time of detection. A practical middle ground would be to aggregate syndromic data by hour rather than by day so that data collected up to the diagnosis time could be used immediately.

Syndromic-surveillance systems are currently deployed and configured to focus on early detection. We conclude that because attack characterization is an equally important role, procedures should be put in place to ensure that the data collected can be appropriately analyzed if an attack occurs and is detected clinically. Such procedures include ensuring that the data are aggregated at a sufficiently fine temporal scale (for example, hourly rather than daily) as well as establishing in advance the statistical measures that should be used to quantify the excess syndromic burden.

Tracking with Cell Phones

We have seen that syndromic surveillance has limited value for bioterrorism defense because the information it provides is nonspecific and because the background noise level of syndromic records makes it difficult to discern the number of records attributable to an attack. With regard to the problem of discerning signal from noise, our simulation results showed that there is a gain in early-detection performance for spatial-temporal surveillance relative to temporal surveillance alone. The spatial-temporal results are based on associating home zip codes with syndromic patients' records. These results could be seen as optimistic because many people are not in their home zip codes during working hours when an attack might occur. From another perspective, though, the results are pessimistic because zip codes comprise relatively large geographic regions. We have developed BACTrack to associate spatially precise, time-resolved geolocation information with syndromic records from a subset of the population [11].

BACTrack would entail recruiting a population of volunteers who agree to have their locations tracked by global positioning system (GPS) technology and to report their general state of health at least once a day through a computerized system. The self-reported health condition will be at the syndromic level—that is, a sick person will be asked to report his or her symptoms and assign them to categories such as respiratory, gastrointestinal, or fever-associated. An algorithm searches for anomalous patterns of illness that are based on current symptom reports combined with the participants' past locations. For instance, the system would take note if an unusually large number of people who had been at a particular bus station at 2 p.m. Tuesday all reported similar, unusual symptoms.

There is typically a delay of several days between exposure to a bio-agent and the initial onset of symptoms, so it is a great benefit to be able to search back in time for clusters of people who currently have similar symptoms. Such a search is not possible with standard syndromic-surveillance systems, because detailed geolocation time histories are not available. Figure 5 illustrates the retrospective search concept for BACTrack, and Figure 6 shows a notional architecture using GPS-enabled cell phones.

In many bioterrorism attack scenarios (e.g., the release of an agent within a building or at an outdoor point), the affected population is concentrated in a small geographic area at the time of the release. However, by

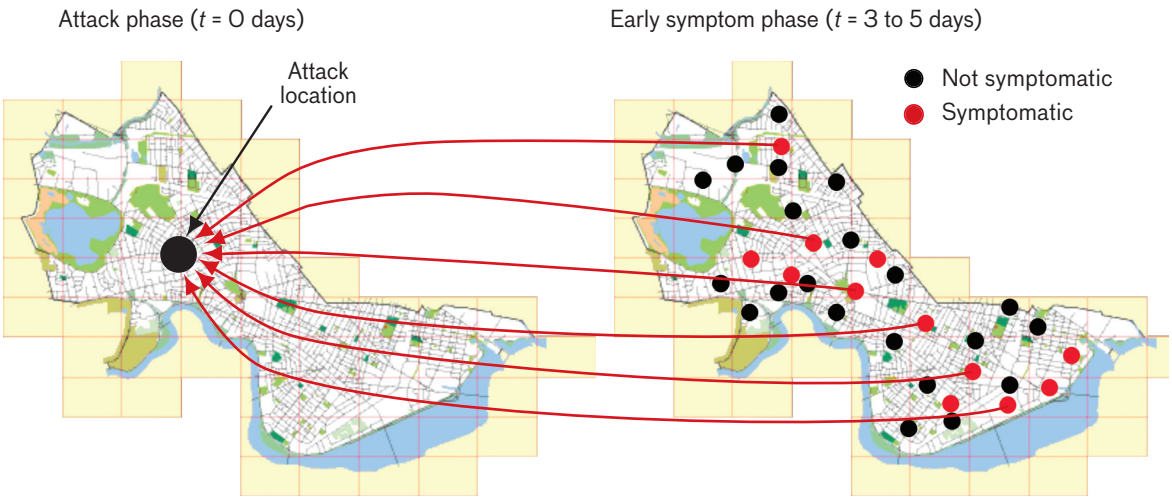

FIGURE 5. In many bioterrorism attack scenarios (e.g., building release or outdoor point release) the affected population is concentrated in a small geographic area at the time of the release. By the time symptoms develop several days later, however, the affected population is likely to be spread out and mixed with the general population. The goal of the Biological-Agent Correlation Tracker (BACTrack) is to detect or characterize an attack by combining current symptom reporting with a historical record of individual geolocations. The system will search for clusters back in time, when people who are symptomatic today were in close geographic proximity; such clusters could reveal the time and location of an attack (shown here in Cambridge, Mass.).

the time symptoms develop several days later, the affected population is likely to be spread out and mixed with the general population. The goal of BACTrack is to detect or characterize an attack by combining current symptom reporting with a historical record of individual geolocations. The system will search for clusters back in time to when people who are symptomatic today were in close geographic proximity, as such clusters could provide clues as to the time and location of an attack.

As is the case with conventional forms of syndromic surveillance, the most promising application of BACTrack would be to characterize an attack that has already been

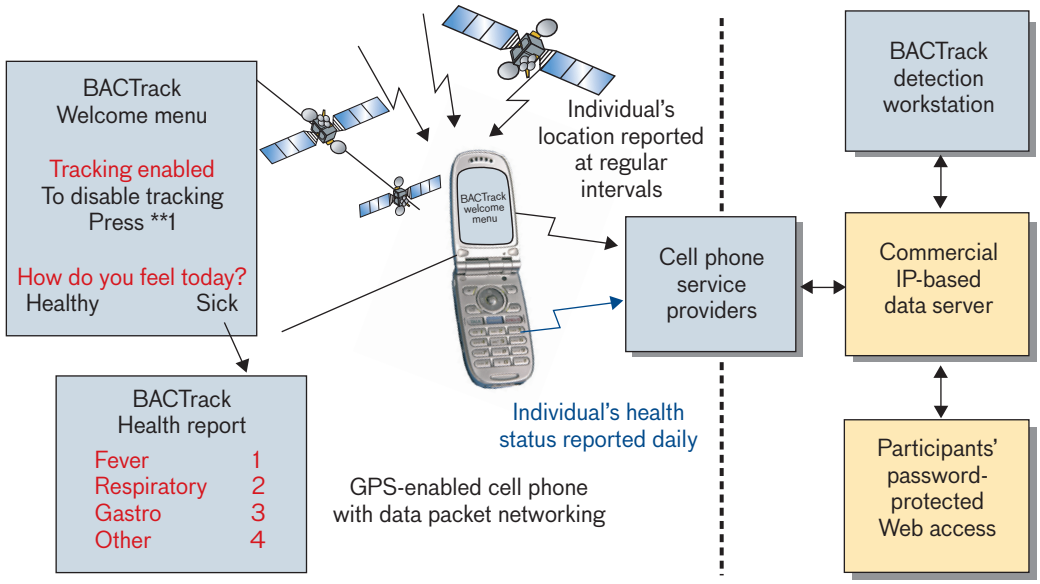

FIGURE 6. BACTrack would rely on the GPS technology found in many of today's cell phones. A specialized application would query the volunteer's health status on a regular basis and report back over an Internet Protocol (IP) network to a central analysis server.

detected through other means. In this sense one could think of BACTrack as an automated system for rapid epidemiology. The standard epidemiological procedure is to interview sick people and attempt to identify a common source of exposure, which can be a time-consuming process. Say, for example, that a sports arena was attacked with anthrax. After an intensive effort spanning several days, investigators would find that many sick patients had been at the same sporting event and would identify it as a likely source of exposure. With a BACTrack system in place, after diagnosis of a single individual, it would be possible to immediately query the BACTrack database for clusters of people showing respiratory symptoms who had recently been in the same place at the same time.

A key question for BACTrack is what percent of the population would need to participate in order for the system to provide useful results. Answering this will require additional modeling and measurement of population patterns. For now, however, we can make some general observations. Even though only a fraction of people are likely to enroll in the system, BACTrack could still be more effective than existing syndromic surveillance in sensing a bioterror attack. BACTrack would make available highly detailed geolocation time histories that could be used to search for exposure clusters based on sick people having been colocated for very short periods of time, such as in a transportation hub. It would also provide for direct symptom reporting without waiting for people to seek medical care and would automatically fuse the symptom reports with geolocation data.

Clearly, recruiting a sufficiently large volunteer population is one of the major challenges in implementing BACTrack. But depending on how events play out, public resistance to participation in such a program could be dramatically reduced. In particular, a large-scale biowarfare attack on the United States would probably cause many people to drop their objections. Thus we recommend preparing to deploy BACTrack when and if the circumstances are such that a large segment of the population will be willing to take part.

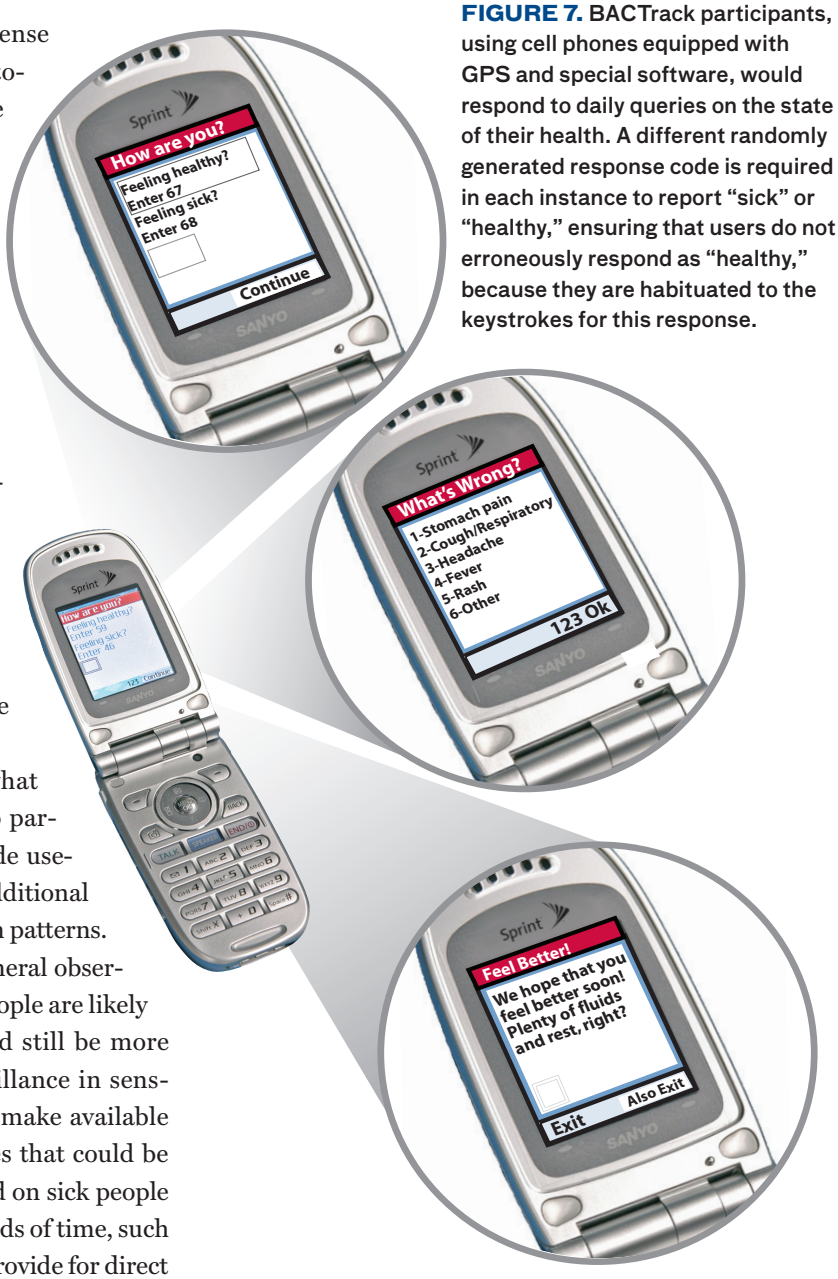

FIGURE 7. BACTrack participants, using cell phones equipped with GPS and special software, would respond to daily queries on the state of their health. A different randomly generated response code is required in each instance to report "sick" or "healthy," ensuring that users do not erroneously respond as "healthy," because they are habituated to the keystrokes for this response.

In cooperation with a national cellular telephone service provider—Sprint—we have developed a BACTrack application to illustrate the data collection functionality on a current-generation cellular handset. The application can be launched remotely from a centralized data server and can query the participants as to their general health status that day. The answers are transmitted back to the central server over the cellular provider's network. Figure 7 shows screen shots of the handset making these health queries. The handset uses its built-in GPS receiver to reg-

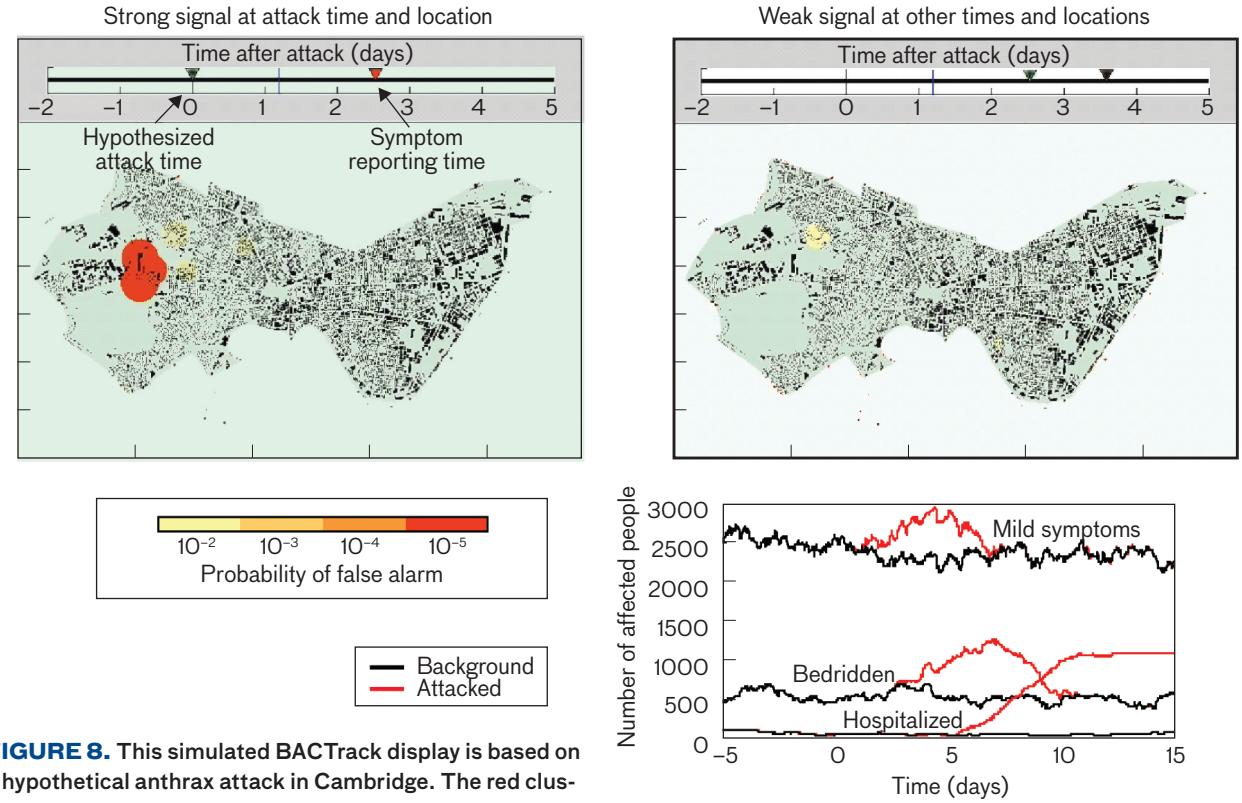

FIGURE 8. This simulated BACTrack display is based on a hypothetical anthrax attack in Cambridge. The red cluster shows an alert with low false-alarm probability at the correct attack location and time (t = 0), based on reported symptoms at a later time (t = 2.5 days). The right-hand side shows that there is not a reliable alert when the system searches for an attack at the wrong time (t = 2.5 days) on the basis of reported symptoms at a later time (t = 3.5 days). The bottom right figure shows the overall health status of the population with and without an attack.

ularly collect geospatial information, which the handset can also transmit to the central server.

We implemented a simulation to demonstrate how BACTrack data could be processed and displayed graphically to highlight candidate exposure clusters. The simulation is set in Cambridge, Mass., and includes movements of 100,000 residents through the course of their daily activities. Figure 8 shows that BACTrack would use symptoms reported two days after a simulated attack to highlight an exposure cluster at the location of the attack. The graphic illustrates how a trained operator could interact with a visualization tool to search for anomalous clusters by varying the hypothesized exposure time as well as the reference time for symptom reports.

This graphical concept of operations is particularly well suited for characterizing an attack after it has been detected by a clinical case finding. The operator would be able to search for evidence of an attack originating at multiple locations in space-time, with likely times and locations showing up as red regions on a map. The red regions could either be targeted for immediate epidemiological follow-up or could be the basis for immediately initiating a mass prophylaxis campaign.

Toward Sensible Surveillance

We have presented the results of two investigations, designed to reduce the time between exposure to a biological agent and distribution of effective treatment to the affected population. Our analyses, and those of others, have shown that the total casualty potential from a bioagent attack increases markedly with increasing time to treatment. Environmental detection networks are scarce and currently inadequate; thus we have focused on time-delay reduction through human-health surveillance. Dedicated investments need to be made in developing rapid-diagnosis technologies such as DNA and protein microarrays for the specific applications of detecting bioagent-induced infectious diseases.

Although syndromic-surveillance systems have been deployed around the country to provide early warning of a biowarfare attack, our analysis suggests that such sys-

tems may have limited value in that capacity. Because the systems have been deployed with the stated goal of early detection, inadequate attention has been paid to configuring them for an application in which they could play an important role: characterizing an attack once it has been detected by clinical diagnosis or by environmental sensor. The requisite data are already being collected; what is needed is to develop and evaluate appropriate algorithms. If an attack occurs and is detected, it will be natural for policy makers to look to the syndromic-surveillance database for an answer to the question of how large the attack is. The operators of these systems should now be starting to take steps to ensure that the most reliable possible answer is available without delay.

The technology required for BACTrack has been demonstrated, and the natural next step is to deploy a small testbed to evaluate its operational feasibility and, perhaps more importantly, the behavior of volunteer participants. Indeed, the greatest challenge for deploying an operational BACTrack system will be recruiting a sufficiently large population of volunteers and ensuring that they continue to be tracked and to reliably report their health status. We do not underestimate the magnitude of this challenge, but we also believe that demonstration of a BACTrack testbed would raise public awareness and could elicit significant enthusiasm to volunteer, at least in certain areas of the country such as Washington, D.C.

In any case, the public mood could change significantly if there were a large-scale biowarfare attack in the next few years. There would then be an increased sense of civic duty and urgency about preparing for another attack. BACTrack could play a very important role by providing a means for the general public to become actively involved in biodefense while also increasing the country's ability to deal with the consequences of a follow-on attack. The government should therefore continue to develop BACTrack by implementing a testbed and integrating the relevant tracking, user interface, and data processing technologies. This development would make it possible to deploy an operational BACTrack system at a future point when there is sufficient public interest to provide a large volunteer population. Such a deployment would dramatically increase the country's ability to rapidly detect and characterize a biowarfare attack, and could save many lives.

Acknowledgements

The topics discussed in this article have had multiple sponsoring authorities. The syndromic surveillance analyses were funded by the Department of Homeland Security (DHS) under programs managed by Teresa Lustig and Steven Buchsbaum. The BACTrack modeling and analyses were supported by funding from the Director of Defense Research and Engineering as well as by DHS under programs managed by Steven Buchsbaum, Keith Ward, and Mike McLaughlin. BACTrack was invented by Larry Candell at Lincoln Laboratory, and its initial development was led by Candell and Ron Hoffeld.

The HSBS study was jointly funded by the Defense Advanced Research Projects Agency (DARPA) and the Defense Threat Reduction Agency (DTRA). It was conducted by a multidisciplinary team from numerous organizations, including the Harvard University School of Public Health, the Naval Research Laboratory, Virginia Polytechnic Institute and State University, Purdue University, Johns Hopkins University, George Washington University, Harvard Medical School, Walter Reed Army Institute of Research, Uniformed Services University of the Health Sciences, Centers for Disease Control, DARPA, DTRA, the Department of Defense Office of General Counsel, and the Army Medical and Research Materiel Command. The principal study author was Darryl Greenwood and topic leaders within Lincoln Laboratory included Bernadette Johnson, Mark Hollis, and Clifford Weinstein. ■

REFERENCES

1. J.A. Jernigan, D.S. Stephens, D.A. Ashford, et al., "Bioterrorism-Related Inhalational Anthrax: The First 10 Cases Reported in the United States," *Emerging Infectious Diseases*, vol. 7, no. 6, Nov.–Dec. 2001, pp. 933–944.
2. C.W. Shepherd, M. Soriano-Gabarro, E.R. Zell, et al., "Antimicrobial Postexposure Prophylaxis for Anthrax: Adverse Events and Adherence," *Emerging Infectious Diseases,* vol. 8, no. 10, Oct. 2002, pp. 1124–1132.
3. D.M. Bell, P.E. Kozarsky, and D.S. Stephens, "Clinical Issues in the Prophylaxis, Diagnosis, and Treatment of Anthrax," *Emerging Infectious Diseases*, vol. 8, no. 2, Feb. 2002, pp. 222–225.
4. W.K. Alexander, Epidemic Outbreak Surveillance (EOS): An Advanced Concept Technology Demonstration, Joint Forces Command and U.S. Air Force Surgeon General.
5. International Society for Disease Surveillance, www.syndromic.org.
6. A. Reingold, "If Syndromic Surveillance is the Answer, What is the Question?" *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, vol. 1, no. 2, 2003.
7. D.L. Buckeridge, D.K. Owens, P. Switzer, J. Frank, and M.A. Musen, "Evaluating Detection of an Inhalational Anthrax Outbreak," *Emerging Infectious Diseases,* vol. 12, no. 12, 2006, www.cdc.gov/ncidod/EID/vol12no12/06-0331.htm.
8. A.A. Szpiro and D.L. Buckeridge, "Beyond Detection: Attack Characterization with Syndromic and Clinical Data," *Advances in Disease Surveillance*, vol. 1, 2006, p. 67.
9. J. Walden and E.H. Kaplan, "Estimating Time and Size of Bioterror Attack," *Emerging Infectious Diseases,* vol. 10, no. 7, 2004, pp. 1202–1205.
10. D.A. Wilkening, "Sverdlovsk Revisited: Modeling Human Inhalation Anthrax," *Proc. Natl. Acad. Sci.*, vol. 103, no. 20, 2006, pp. 7589–7594.
11. R. Hoffeld and L. Candell, "Technique for Rapid Detection and Localization of Attacks with Biologic Agents," *Syndromic Surveillance: Reports from a National Conference*, 2003, *Morbidity and Mortality Weekly Report,* 2004, www.cdc.gov/MMWR/preview/mmwrhtml/su5301a55.htm.

ABOUT THE AUTHORS

Adam Szpiro is a senior fellow in the Department of Biostatistics at the University of Washington, where his research focuses on statistical methods in air pollution epidemiology. He is a former staff member in the Biodefense Systems group at Lincoln Laboratory, where he assessed requirements for sensing and responding to an urban biological warfare attack; he also initiated the Lincoln Laboratory infectious model for disease epidemics. He has a doctorate in applied mathematics from Brown University.

Bernadette Johnson is leader of the Biodefense Systems group, where she directs work on sensor development, laboratory and field measurements, biodetection forensic techniques, and systems analyses for military and civilian biodefense. From 1993 through 1996, she directed the Environmental Monitoring Project, which aimed to adapt Lincoln Laboratory technologies to environmental-monitoring applications. She has a doctorate in plasma physics from Dartmouth College.

David Buckeridge is an assistant professor of epidemiology and biostatistics at McGill University in Montreal, where he specializes in the informatics of public health surveillance. He has an M.D. from Queen's University in Canada, a master's degree in epidemiology from the University of Toronto, and a doctorate in biomedical informatics from Stanford.