Chapter Three

MULTISENSOR SYSTEM TO IMPROVE MINE DETECTION CAPABILITY

No single mine sensor has the potential to increase the probability of detection and decrease false alarm rates for all types of mines under the wide variety of environmental conditions in which mines exist. As is clear from the previous chapter, each innovative method is subject to limitations under certain conditions of environment and mine type. Thus it is unlikely that any one technology will provide the breakthrough necessary to substantially improve humanitarian demining operation times.

To achieve substantial decreases in mine detection time while maintaining high probabilities of detection, a paradigm shift is needed in mine detection R&D. Rather than focusing on individual technologies operating in isolation, mine detection R&D should emphasize the design from first principles and subsequent development of an integrated, multisensor system that would overcome the limitations of any single-sensor technology. The goal of such R&D should be to produce a single system comprised of many different types of sensors and algorithms for combining the feedback from all sensors in an optimal manner.

Multisensor systems that combine technologies with different sources of false alarms could substantially decrease the false alarm rate. As an example, Table 3.1 lists the individual sensor technologies that, based on the evaluations in the previous chapter, appear most promising for close-in mine detection for humanitarian operations. As shown, the methods have substantially different sources of false

Table 3.1

Sources of False Alarms for Selected Mine Detection Technologies

Detection Technology	Primary Source of False Alarms
EMI	Metal scrap, natural soil conductivity, and magnetism variation
GPR	Natural clutter (roots, rocks, water pockets, etc.)
Acoustic/seismic	Hollow, man-made objects (e.g., soda cans)
Fluorescent polymers	Explosive residues
NQR	Radio frequency interference

alarms. Multisensor systems also could increase the probability of finding diverse types of mines and of operating effectively in a range of environments. Table 3.2 shows the most promising technologies and the mine and environmental conditions for which they are best suited. As can be concluded from Tables 3.1 and 3.2, combining technologies could yield a system that is robust across diverse environments and mine types.

The potential of multisensor systems to improve detector operating characteristics has been demonstrated empirically. In a comparative field study of different sensor fusion algorithms, Gunatilaka and Baertlein (2001) showed that fusion of EMI, GPR, and infrared data achieved better than a factor of 8 improvement in the probability of false alarm compared with the best individual sensor, at PD = 1. Similarly, Collins (2000) found that advanced fusion of EMI, GPR, and magnetometry signals achieved between a factor of 5 and 12 improvement in probability of false alarm compared with simple thresholding on raw data of individual sensors, at PD = 0.8. Finally, in an extensive field test of five vehicle-mounted systems consisting of metal detectors, GPR, and infrared sensors conducted by the Institute for Defense Analyses (IDA), fusion of the three signals improved operating characteristics for all systems compared with any single sensor or pair of sensors (Rotondo et al., 1998). These studies, and others cited in the appendixes, demonstrate that a well-designed multisensor system has the potential to substantially improve performance relative to detectors that use a single-sensor technology.

Table 3.2

Mine Types and Soil Moisture Conditions for Which Selected Mine Detection Technologies Are Best Suited

	Mine Types and Soils for Which the Technology Is Most Effective					
Detection Technology	Metal Mines	Low-Metal Mines	TNT Mines	RDX Mines	Dry Soil	Wet Soil
EMI	√		√	√	√	√
GPR	√	√	√	√	√	
Acoustic/seismic	√	√	√	√	√	√
Fluorescent polymers	√	√	√	√		√
NQR		√		√	√	√

KEY DESIGN CONSIDERATIONS

A multisensor system that achieves performance superior to a single-sensor system across a diversity of environments and target types will require effective and flexible methods for combining information from multiple sensors of different modalities. Multisensor systems can combine or “fuse” the information from the component sensors in a variety of ways, characterized by the stage of signal processing at which the information is combined. Broadly speaking, the different fusion methods can be categorized as “decision-level,” “feature-level,” or “data-level” (Ackenhusen et al., 2001; Gunatilaka and Baertlein, 2001; Dasarathy, 1994). In decision-level fusion, each component sensor of the system provides the operator with a declaration decision from independently processed signals, and these are combined to make the overall declaration decision. For example, HSTAMIDS (see Chapter Two) produces two separate signals—one from the GPR detector and another from the EMI detector—and the operator must decide whether to investigate an item, depending on the signals received. “Hard” decision-level fusion bases the overall decision on only the individual binary decisions (declare/non-declare) of the component sensors, generally using simple rules (e.g., Boolean “and,” “or,” or majority voting). Alternatively, if individual declaration decisions are augmented with some measure of confidence, “soft” decision rules that give more weight to more reliable decisions are possible. In either case, the overall decision is based on only the independently processed signals from the individual sensors. This is in contrast to feature- or data-level fusion, in which the signals generated by each sensor are combined algorithmically to present the operator with a single signal on which to base the declaration decision. That is, rather than combining the processed outputs of various sensors, these lower-level fusion methods jointly process the received physical information at the signal level. Data-level fusion combines the raw data collected by each sensor, while feature-level fusion combines information about informative "features" extracted from the raw signals.

Situations occur in which one type of fusion might be advantageous, relative to another. For example, if the multiple sensors detect the same type of target but have different confounders, decision-level fusion helps to reduce errors and improve operating characteristics. However, feature-level fusion would likely be more effective for an array of sensors, each of which is designed to detect a different kind of target. Moreover, there are theoretical reasons to believe that properly implemented fusion at lower information levels (e.g., data or feature) should be superior to fusion at higher (e.g., decision) levels. This is because such fusion can make more efficient use of the available information by exploiting the simultaneous characteristics of the signals ignored by decision-level fusion. Lower-level fusion also in principle offers a pragmatic advantage because it operates as an integrated unit, presenting the user with a single signal in a manner similar to traditional EMI detectors. "Primary" and "confirmatory" sensing would be transparent to the operator. Alternatively, when multiple signals are presented to the operator, multiple operating thresholds must be chosen. This makes performance optimization more cumbersome.

Most work to date has used decision-level fusion, which has been shown to reduce false alarm rates relative to individual sensor performance while maintaining high probabilities of detection. Data- or feature-level fusion is more difficult, less mature, and thus far has been shown to offer relatively modest incremental improvements to decision-level fusion. For example, in a performance study of decision- versus feature-level signal processing on data collected from EMI, GPR, and infrared sensors, Gunatilaka and Baertlein (2001) demonstrated that hard decision-level fusion did not perform appreciably better than the best of the individual sensors (GPR in this case). However, soft decision- and feature-level fusion offered marked benefits, with the feature-level algorithm providing a modest but meaningful additional benefit relative to the soft decision-level methods.

Because feature- and data-level fusion are relatively new research areas, it is uncertain how rapidly and to what degree their theoretical advantages over decision-level fusion will result in practical applications. Thus, in designing the multisensor system we propose, we believe it is important to keep an open mind about how to process information from multiple sensors. Researchers should be receptive to combining information at whatever level is necessary to achieve optimal performance in practice. However, given the potential to improve operating characteristics and simplify multisensor system operation (by presenting a single signal), research should continue on advanced lower-level fusion algorithms.

Another formidable challenge to the next-generation multisensor system is having the flexibility to be effective across a broad range of field conditions and target types. Regardless of the level at which information from the component sensors is combined, performance using a fixed configuration of operating thresholds will vary tremendously across the diversity of field conditions encountered in humanitarian demining operations. It is impossible to determine optimal thresholds for all possible field conditions a priori. Thus, achieving optimal performance will require adjustments to operating thresholds on a case-by-case basis. Such adaptation may not be practical for some military demining operations but is practical for most humanitarian demining operations. The mode of adaptation conceptually falls along a continuum of end-user responsibility. At one extreme, the user is completely responsible for optimizing the system for the field, with no built-in system intelligence for this task. At the other extreme, the system is entirely self-calibrating and learns about the field in an automatic, real-time fashion. While the latter ideal may be a worthy long-term goal, it is unrealistic in the short term, and initial development efforts should compromise between system and user optimization. Through both field testing and aggressive environmental and target-sensor interaction modeling, it should be possible to learn what environments and targets are most challenging to the system and how the system can be altered to handle them effectively. This information can be embedded into the system logic, and the user would then input information about a small number of high-leverage field condition variables that would result in gross alterations to system thresholds. The remaining user performance tuning could then be along some manageably simple dimensions, ideally of complexity on par with current EMI detectors. We envision that, as research progresses, more of the local adaptation can be off-loaded to the system.

POTENTIAL FOR A MULTISENSOR SYSTEM TO INCREASE MINE CLEARANCE RATE

The primary motivation for pursuing multisensor systems is to reduce the false alarm rate and in turn decrease the amount of time spent on mine detection. In requesting this study, OSTP asked that RAND S&TPI assess whether improved detection methods could increase the speed at which mines can be cleared by an order of magnitude. While we conclude that multisensor systems are the most promising path for improved performance, we also determine that it is unlikely that such systems will achieve order-of-magnitude improvements in clearance rates in the near future. There are two primary reasons for this conclusion. First, reductions in the false alarm rate do not bring about equal reductions in mine clearance time, even under the assumption that the time spent removing actual mines is negligible. Depending on the specific environment, a nontrivial portion of the total clearance time may be spent on such "overhead" operations as vegetation clearance. In some cases, these preclearance activities may require so much time that even a detection system with perfect operating characteristics could not achieve an order-of-magnitude improvement in clearance rates. For example, Treveylan estimates that an improved detection system that eliminated 99 percent of false alarms would speed overall clearance rates by 60–300 percent of the original clearance rate (meaning 1.6–4.0 mines could be cleared in the time now required to clear one mine), depending on the density of vegetation (Treveylan, 2002a,b). This is in contrast to the hundredfold (i.e., 9,900-percent incremental) improvement that is theoretically possible if only the time spent investigating false alarms is considered.

The second reason we conclude that order-of-magnitude improvements are unlikely in the near future stems from tradeoff between decreasing the false alarm rate and increasing the probability of detection. For single-sensor systems, reducing the false alarm rate decreases the probability of detection. For the near future, the same phenomenon is likely to occur even for a well-designed multisensor system. Because of the resulting increased residual risk from more missed mines, this is not acceptable. Thus, the goal of minimizing the false alarm rate will continue to be limited by the need to maximize the probability of detection.

Although order-of-magnitude decreases in mine clearance rates are unlikely to be achieved in the near term, focused investments in multisensor system development would yield substantial savings in time and money. Reductions in false alarms translate into reductions in operation times, even if these reductions are not proportional; therefore, operation costs should decrease substantially. For example, the United Nations has estimated that the average cost of clearing a mine is $300–1,000 per mine (Hubert, 1998). Based on this average, the cost to clear all 45–50 million mines worldwide will total $14–50 billion. Given that a large fraction of these costs derive from paying deminers for time spent investigating false alarms, even moderate time savings could save billions of dollars. For example, Treveylan (2002a,b) has estimated that if 50 percent of false alarms were correctly declared as nonhazardous, then demining speed would improve by approximately 30 percent in a highly vegetated country, such as Cambodia, and by approximately 60 percent in a minimally vegetated country, such as Afghanistan, compared with a detector that has no ability to distinguish false alarms from mines. If one assumes that the improved system would increase demining speed by 30–60 percent worldwide and that such gains in speed translate directly into cost reductions, then the improved detector would save 23–38 percent of the total demining cost of $14–50 billion. Improvements beyond an ability to correctly identify 50 percent of false alarms are expected, with proportionately higher decreases in mine clearance time and costs. For example, Treveylan (2002a,b) estimates that a system that correctly identified 90 percent of false alarms would speed demining by 60 percent in Cambodia and 200 percent in Afghanistan.

In addition, the development of an effective multisensor system, and the resulting reduction in false alarms, has numerous tangible benefits beyond reduction in operation times. The vast number of false alarms encountered in demining operations may foster inattention and carelessness, which may increase the occurrences of demining accidents. A reduction in the false alarm rate thus has the potential to improve demining safety. Moreover, although the discussion of multisensor systems has focused on the potential for minimizing the false alarm rate, it is possible that modest false alarm rate reductions could be coupled with improvements, rather than degradations, to the probability of detection. This ultimately improves public safety. Also, the R&D necessary to pursue the integrated multisensor system will require aggressive interdisciplinary efforts by top researchers and is likely to result in technological advances that can be leveraged in other applications of public interest (such as chemical weapon detection, airline safety, drug enforcement, military countermine operations, and improvised explosive device detection). Finally, even if order-of-magnitude improvements are not achievable in the near term, incremental improvements may lead to an order-of-magnitude improvement over time.

CURRENT U.S. R&D INVESTMENT IN MINE DETECTION TECHNOLOGIES

The primary conclusion of this report is that development of an advanced, integrated multisensor system is the most likely pathway to substantial improvements in humanitarian demining operations. However, the federal government currently has no research program that is directed specifically at designing such a system. While the United States funds research related to humanitarian mine detection, the amount allocated is very small. Further, existing funding is not optimized toward design from first principles of an advanced multisensor system.

Currently, federal funding for R&D on humanitarian mine detection is limited. The only federal R&D program dedicated to humanitarian demining is the Department of Defense Humanitarian Demining R&D Program (HD program), established in 1995 and administered by the Assistant Secretary of Defense for Special Operations/Low-Intensity Conflict. Total funding for this program in 2002 was $13.5

million (UXO Center of Excellence, 2002). Of that amount, $4.9 million was allocated for detection technologies.[1] Nearly half ($2.2 million) of the detection technology budget was spent on wide-area detection R&D (for use in remote identification of potential mine field locations). Thus the total amount available for R&D on close-in detection technologies was $2.7 million.

Existing federal funding for humanitarian mine detection is focused primarily on optimizing the performance of EMI and GPR systems. As shown in Figure 3.1, most of the $2.7 million spent on close-in methods in 2002 was concentrated on GPR and EMI. The only other detection method included was chemical vapor–sensing technologies. No funding was invested in researching a multisensor system.

The HD program emphasizes traditional detection methods because it is an applied research program and is by design not positioned to

RAND*MR1608-3.1*

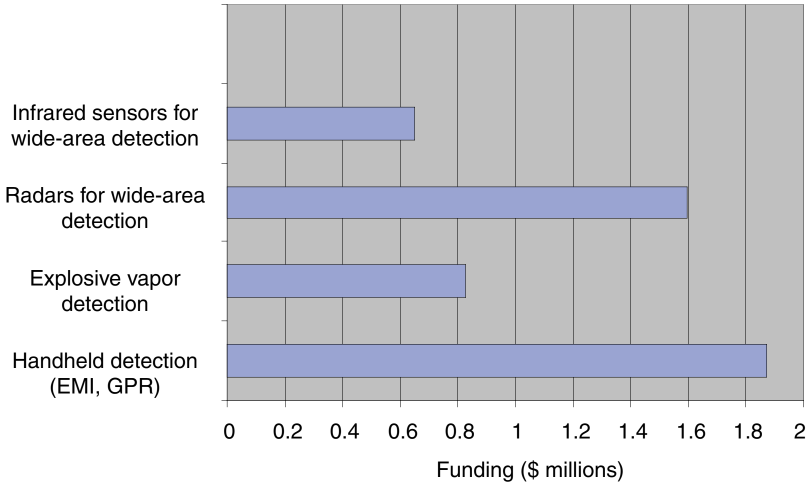

NOTE: Total funding for detection technologies was $4.9 million. Data in figure are from R. Weaver and S. Burke (personal communication, August 15, 2002).

Figure 3.1—Detection Technologies Funded by the U.S. Humanitarian Demining R&D Program in Fiscal Year 2002

---

[1] Information obtained from personal communication with R. Weaver and S. Burke, Army Night Vision and Electronic Sensors Directorate, Ft. Belvoir, Va., August 13, 2002.

conduct basic research. The U.S. Department of Defense organizes its research into five tracks corresponding to level of maturity of the technology, with track 6.1 dedicated to basic research projects and tracks 6.2–6.5 dedicated to applied research and engineering design leading to an operational system. The HD program is strictly an applied research program, focused on development (corresponding to the 6.2 and 6.3 R&D levels) of concepts for which the critical basic research questions already have been answered. The existing funding allocation makes sense given the program's applied nature. However, the knowledge necessary for engineering development of a multimodal sensor of the type recommended in this report does not exist and would need to be developed through basic research.

The Army Research Office funded a five-year basic research effort through the Department of Defense Multidisciplinary University Research Initiative (MURI) from 1996 to 2001. It provided a total of $3.2 million per year to three university consortia for basic research on a variety of mine detection technologies. Some of this funding was for development of theory to support fusion of information from multiple sensors. However, a comprehensive, predictive theory for fusing signals or information from multiple landmine sensors still is not developed (J. Harvey, personal communication, 2002). Like other MURIs, this initiative covered only a five-year period and is not being continued. Current Army Research Office MURI investments, initiated in 2002, are funding three consortia a total of $9.2 million over the next five years (about $1.8 million per year) to explore three areas related to mine detection: (1) real-time explosive-specific chemical sensors, (2) the science of land target spectral signatures, and (3) detection and classification algorithms for multimodal inverse problems (UXO Center of Excellence, 2002). While some of the results of these efforts may be applied to mine detection, mine detection is not the exclusive focus of these new MURI projects.

The majority of U.S. funding for mine detection research is allocated to development of systems for countermine warfare. In 2002, the federal government invested $106 million in countermine research, about $75 million of which was spent on detection technologies (UXO Center of Excellence, 2002). The humanitarian demining community could leverage the technology developed for countermine operations, although some of the requirements are significantly different. (For example, detectors for countermine operations need not achieve the near-perfect probabilities of detection required for humanitarian demining.) The distribution of funds for countermine research resembles that for humanitarian mine detection in that the funding is concentrated on traditional GPR and EMI systems because of the emphasis on applied research.

As in the humanitarian demining program, the knowledge needed to field a multisensor system that yields one signal has not been developed through the countermine program. As discussed earlier, the development of HSTAMIDS, which combines GPR and EMI, has been a major focus of countermine research. However, because HSTAMIDS was developed in response to the immediate need for a system that works better than traditional metal detectors, it does not optimize use of the signals from the individual sensors. It combines separately designed EMI and GPR technologies on a single platform but does not necessarily represent the optimal combination of these two systems. Further, the output it produces is not aided by multisensor decision algorithms. The operator hears two separate signals from the EMI and GPR sensors and is not assisted in deciding which combination is most likely to indicate the presence of a mine. The decision about whether to declare an item a mine is not straightforward.

In some cases, strong signals from both the EMI and GPR detectors signal a mine, but in other cases the absence of a signal from one in conjunction with a strong signal from the other would indicate that a mine may be present. For example, consider a low-metal mine buried in dry sand with wet spots in the sand. An EMI detector will sense the currents induced in the metal or in the wet spots as being "targets." A GPR detector will identify a strong return signal from the wet sand and a relatively weak return signal from the plastic mine case, which has a similar dielectric constant as dry sand. In this case, the presence of a strong EMI response and the absence of a GPR response would indicate the presence of a target. This example illustrates the need for sensor fusion algorithms to assist the deminer in making declaration decisions. HSTAMIDS does not include such algorithms.

In sum, currently there is no basic research program focused on developing the fundamental knowledge needed to engineer novel multisensor detection systems for humanitarian demining. The funding allocated for humanitarian mine detection R&D is limited, and it is not dedicated to research toward a multisensor system. R&D conducted under countermine programs is well funded but is not optimized for multisensor system development. No basic research program exists to continue development of the theory needed to support a multisensor system.

RECOMMENDED PROGRAM FOR PRODUCING AN ADVANCED MULTISENSOR SYSTEM

Development of an advanced multisensor mine detection system will require new, targeted R&D funds. The program should focus on the following areas:

1. algorithms for fusion of data from the individual sensors (to develop the theory necessary to support an advanced multisensor system);
2. integration of component systems (to address engineering issues associated with combining multiple sensors as part of a single device);
3. detection of the chemical components of explosives (to further develop components of the multisensor system that would search for explosives rather than for the mine casing and mechanical components); and
4. understanding of the fundamental physics of how the soil conditions in the shallow subsurface environment affect different sensors (to allow the development of models to predict performance across a range of environments).

This should be a long-term program, continued until the basic scientific and engineering knowledge needed to field an integrated, multisensor system is complete.

Because the long-term program will require a sustained commitment of substantial resources, we recommend that a short, relatively inexpensive preliminary study be conducted. The study should not focus on new research; rather, it should consolidate the existing empirical and theoretical work on sensor fusion and signal processing, with a focus on applications in landmine detection. The results of such a comprehensive literature review and assessment will help to narrow the focus of the long-term program we recommend to the most-promising avenues. This report covers some of the most-promising results to date; however, signal processing and sensor fusion are rapidly growing research areas and are being pursued by a broad array of academic, industrial, and governmental organizations in a diversity of disciplines. The resulting fragmentation of key empirical results and theoretical investigations makes it difficult to distinguish the rhetoric of sensor fusion and signal processing from the reality. Prior to investing in the long-term program, we believe it is crucial for decisionmakers to have a comprehensive assessment of the potential benefits and limitations of signal processing and sensor fusion for landmine detection, expert-informed guidance about which combinations of technologies and algorithms are most promising, and realistic estimates of the potential for improved operating characteristics.

COST OF DEVELOPING A MULTISENSOR SYSTEM

As in any R&D initiative, the costs of developing a multisensor system are difficult to predict in advance. If the path forward for developing a next-generation detector were clearly defined, this would no longer be a research issue but strictly an engineering problem. Estimates of total anticipated research costs can be made based on prior experience, but it is important to keep in mind that these estimates may change over time as additional knowledge is gained.

To predict what investment might be needed to develop the next-generation mine detector, we used actual R&D costs of HSTAMIDS as a model. Table 3.3 shows actual costs of the various stages of HSTAMIDS development in the third column. The fifth column shows predicted costs of developing the next-generation multisensor system (excluding the costs of the preliminary study to consolidate existing research). The predictions for the next-generation detector are based on two assumptions:

1. that basic research to develop the new system will require three times as many researchers as participated in HSTAMIDS development, and

Table 3.3

Estimated Costs and Time Required to Develop a Next-Generation Multisensor Landmine Detector

R&D Stage	Time Required for HSTAMIDS	Cost for HSTAMIDS	Estimated Time for Next-Generation Multisensor System	Estimated Cost for Next-Generation Multisensor System
Basic Research	4 years (1990–1994)	$5 million	5–8 years	$50 million
Prototype Development	2 years (1994–1996)	$8 million	2 years	$10 million
Demonstration and Validation	5 years (1996–2001)	$33 million	5 years	$40 million
Engineering and Manufacturing Development	4 years (2001–2005)	$27 million	4 years	$35 million

2. that once the basic research is completed, the remaining stages of engineering the new system should cost approximately the same as the equivalent stages of building HSTAMIDS, adjusted by 1.1 percent per year for inflation.

The basic research for the next-generation detector will be considerably more complicated than that required for HSTAMIDS for two primary reasons. First, HSTAMIDS incorporates two mature technologies (EMI and GPR), whereas the next-generation system might include newer detection methods, such as NQR and/or acoustic/seismic sensors. Performance of these newer methods cannot be optimized or predicted without additional research. Second, HSTAMIDS does not include advanced algorithms for sensor fusion. The theoretical basis needed to accomplish advanced sensor fusion is not yet complete. As Table 3.3 shows, we expect that eight years of basic research will be required to support an advanced, multisensor detector, but it is possible that this research could be compressed into five years, with the total amount of funding remaining the same but with a higher amount spent each year. A prototype system could be available within two years of completing the basic research.

Table 3.4 summarizes the types of basic research required to support development of a next-generation, multisensor mine detector, as

Table 3.4

Itemization of Basic Research Requirements and Costs over the Next Five Years for a Next-Generation Multisensor Mine Detector

Research Area	Estimated Number of Researcher Years over the Next Five Years	Estimated Cost over the Next Five Years	Anticipated Results After Five Years of Research
Algorithms for sensor fusion	40	$10.00 million	Minimal set of sensor-level fusion algorithms for specific sensor suite
Integration of component sensor technologies	25	$6.25 million	Multisensor prototype detector with three to four sensor technologies
Explosives detection technologies	50	$12.50 million	Set of sensors suitable for integration in multisensor prototype; three to five would be candidates for immediate integration, with remainder used for backups or requiring long-term research for midlife upgrades
Subsurface environment effects on sensors	10	$2.50 million	Understanding of major soil parameters that affect the sensors identified above; set of simple tests that can be performed in situ to provide information to improve or predict performance of those sensors

recommended in this report. Research costs are based on the judgment and experience of the S&TPI/OSTP Landmine Detection Technology Task Force. The amounts shown here assume that the basic research work is not compressed—i.e., that significant progress will occur after five years, but that funding is not sufficient to complete the work in this time frame.

As shown in Tables 3.3 and 3.4, the total estimated costs for basic research toward developing a next-generation detector are $50 million over five to eight years, or $6.25–10.00 million per year. (The higher annual figure assumes that the research would be compressed into five years.) The total anticipated costs, including basic research, prototype development, demonstration and validation, and engineering and manufacturing development are \$135 million over 16–19 years (\$7.1–8.4 million per year). As shown, a prototype system could be available within seven years at a total cost of \$60 million (\$8.6 million per year) if the basic research were compressed with higher up-front spending.

SUMMARY OF RECOMMENDATIONS

In sum, no basic research program focused on developing the fundamental knowledge needed to engineer novel multisensor detection systems for humanitarian demining currently exists. We recommend that the United States invest in developing such a system to enable more rapid, safer clearance of antipersonnel mines. Such a system would save billions of dollars in the cost of mine clearance, and the resulting advances in signal processing and sensor fusion would be transferable to many other disciplines and applications.

The first phase of the multisensor mine detection program should be a limited, proof-of-concept study (costing less than \$1 million) that would consolidate existing research on sensor fusion and signal processing. This study would identify the most-promising directions for multisensor development and the key information gaps.

Once the proof-of-concept study is completed, the full multisensor development initiative would begin with research focused on the following broad areas:

- algorithmic fusion of data from individual sensors (\$2.0–3.2 million per year);
- integration of component sensor technologies (\$1.25–2.00 million per year);
- detection of chemical components of explosives (\$2.5–4.0 million per year); and
- modeling of the effects of soil conditions in the shallow subsurface on sensor technologies (\$0.5–0.8 million per year).