SAND2009-1991C

An End-to-End Approach to Developing Biological and Chemical Detector Requirements

Nerayo P. Teclemariam*, Liston K. Purvis, Greg W. Foltz, Todd West, Donna M. Edwards, Julia A. Fruetel, and Nathaniel J. Gleason
Sandia National Laboratories, P.O. Box 969, Livermore, CA 94551-0969

ABSTRACT

Effective defense against chemical and biological threats requires an "end-to-end" strategy that encompasses the entire problem space, from threat assessment and target hardening to response planning and recovery. A key element of the strategy is the definition of appropriate system requirements for surveillance and detection of threat agents. Our end-to-end approach to venue chem/bio defense is captured in the Facilities Weapons of Mass Destruction Decision Analysis Capability (FacDAC), an integrated system-of-systems toolset that can be used to generate requirements across all stages of detector development. For example, in the early stage of detector development the approach can be used to develop performance targets (e.g., sensitivity, selectivity, false positive rate) to provide guidance on what technologies to pursue. In the development phase, after a detector technology has been selected, the approach can aid in determining performance trade-offs and down-selection of competing technologies. During the application stage, the approach can be employed to design optimal defensive architectures that make the best use of available technology to maximize system performance. This presentation will discuss the end-to-end approach to defining detector requirements and demonstrate the capabilities of the FacDAC toolset using examples from a number of studies for the Department of Homeland Security.

Keywords: Biological, chemical, terrorism, requirements, analysis, countermeasures, detection, defense, simulation

1. INTRODUCTION

As recently highlighted by the Commission on the Prevention of Weapons of Mass Destruction Proliferation and Terrorism report[1], the threat of terrorism using weapons of mass destruction (WMD), such as biological or chemical weapons, remains a credible threat even though the United States has not experienced an attack since the anthrax letter mailings of 2001. Effective defense against chemical and biological threats requires a comprehensive "end-to-end" strategy that considers the entire problem space, from surveillance and detection to threat assessment, target hardening, response planning, and recovery. The approach begins by understanding the likely threats that the defensive system will be designed to protect against and the probable targets of those threats. Target vulnerabilities are assessed, and hardening actions and countermeasures are evaluated for their ability to reduce the likelihood or impact of an attack. Detection and surveillance systems are designed to provide information to trigger countermeasures. The deployment and design of these defense architectures is optimized to ensure broad system coverage and the delivery of actionable information. Finally, restoration and recovery strategies assist in returning the target to normal operations as quickly as possible after an attack.

Significant effort and focus has been placed on the creation of detection and monitoring systems to inform or enable protection measures that reduce and prevent exposure following an attack[2,3]. As investments continue to be made in the development and deployment of future detection technologies, system requirements must be established with the entire end-to-end scope in mind to ensure that systems are adequately addressing mission objectives.

This paper provides an overview and illustrates the benefit of using an end-to-end approach to defining chemical and biological detector requirements for all phases of the detector development process. The discussion will begin by introducing the Facilities Weapons of Mass Destruction Decision Analysis Capability (FacDAC) toolset developed by

* nptecle@sandia.gov; phone 1 925 294-4823; fax 1 925 294-3870; www.sandia.gov

Sandia National Laboratories to model chemical and biological defense systems. Next, Sandia's end-to-end approach to evaluate detector systems and develop requirements will be discussed in detail. Finally, the approach's applicability will be demonstrated in various stages of the detector development process including guidance for early-stage technology investment, development-stage system evaluation, and application and deployment of assets.

2. FACDAC CHEM/BIO VENUE PROTECTION MODELING TOOLSET

FacDAC is an integrated system-of-systems toolset that enables the simulation of indoor and outdoor chemical and biological incidents and responses. Modules featured in the toolset include facility geometry, HVAC and indoor airflow models, simulated attacks, population movement, exposure and disease progression models, outdoor dispersion models, public health care network, facility response options, signal interpretation, and detection systems. Coupling of the individual modules allows for biological and chemical attacks and responses to be characterized from initial release to resolution. The primary components of FacDAC include models that capture airflow movement within and outside facilities, contamination transport, population movement throughout a building, and toxic chemical and biological agent health effects (e.g., dose-response curves and disease progression models). The airflow models provide accurate representation of air movements through a building. Contaminant transport models account for the dispersion, filtration by air handling systems, deposition via gravitational settling, and other loss processes involved during a release. Population movement models capture the passage of people throughout a facility. To account for variations in paths taken by individuals through venues, a probabilistic approach is used. Each person entering a facility was assigned a unique schedule that could include points of entry and exit, and time waiting in designated areas (e.g., ticketing, security checkpoints, and boarding areas). Moreover, the dimensions of doorways and stairs were explicitly included in the population models as they can create resistance to movements of large populations (i.e., bottlenecks). Health effects models characterize the impact of toxic chemicals and biological agents on a diverse population. The amount of chemical or biological agent exposure that each person receives was determined by the release location, release size, release time, agent transport through the facility, and each person's movement through the facility.

To account for variations in weather and building operations, a Monte Carlo approach is used to build a library of building "states" that encompass a broad range of operating conditions.. Parameters such as the outdoor temperature and wind direction, the efficiency of filtration in the heating, ventilation, and air conditioning (HVAC) system, and whether or not doors are open or closed within a facility will have significant impacts on air flows. By placing probability distributions on these parameters and running simulations, a database of hundreds of thousands of scenarios is generated for use in analyses. FacDAC analyses consider these variations to ensure requirements are robust and not dependent on a single set of operating conditions.

3. APPROACH TO DEVELOP DETECTOR REQUIREMENTS

When defining detector requirements, it is important to choose an appropriate metric for evaluating performance and to quantify whether a architecture satisfies the overall objectives of the program. Types of metrics employed to evaluate chemical and biological detection system deployments vary; however, two categories of metrics are frequently used: impact-based and mass-based. Impact-based metrics capture the outcomes and consequences of attacks (e.g., number of infected people), while mass-based metrics characterize the size of release (e.g., one kilogram). . A main advantage of using impact-based metrics to evaluate detector requirements and performance is that focus is on mass releases that have significant impact and can also be mitigated (i.e., mass releases that do not have significantly impact will be deemphasized). In the work discussed in this paper, the impact-based metrics used include the number of people who would be intoxicated (for chemicals) or infected (for biological agents) and the "impact reduction," which is the difference in the number of casualties resulting from detection and response compared to no detection. The number of people infected or intoxicated can be determined from standard dose-response curves that relate agent exposure to health effects (e.g., infected or incapacitated). The "impact reduction" metric is especially useful for evaluating early warning detection systems that must be coupled with an effective response to protect populations. This metric is computed by

first determining the number of casualties that would result from a specific attack if no external response (e.g., detector alarm) was initiated. It is important to note that self-evacuation and people-triggered evacuation, due to reaction to chemical odor or the observation of incapacitated people, are explicitly included in the analysis. The difference in the number of casualties without detection and with detection plus response is then determined and used to characterize the system.

The process of detector development contains three major stages: early, development, and application. The early stage of detector development involves the definition of specific performance targets needed to satisfy the overall objective of the detection system and the establishment of guidance on potential technologies to pursue for incorporation in the system. Performance targets include the sensitivity of the detector to identify agents of interest, selectivity of the detector to discern target agents from background, and false positive rates. In the development stage, detector technologies have been selected and their performance is evaluated. Competing technologies are compared and performance trade-off analyses are conducted to determine, which technology is most applicable to addressing the system goals. The application stage involves the deployment of developed detection systems into their operational environment. The end-to-end approach to developing biological and chemical detector requirements will be explained by walking through the process in detail for the first stage of the detector development process.

A schematic view of the major components involved in the process of recommending which detection technologies to develop is featured in Figure 1. To appropriately design requirements for detection systems one must have a clear understanding of the threat the detection system is designed to mitigate and the response actions that will be triggered by the system. Aspects of the threat can be compiled into a Design Basis Threat (DBT), or a set of credible scenarios that span a representative threat space over which the detector architecture will be expected to function. Understanding the response environment will also assist in defining requirements by identifying types of information the system should produce to trigger a response. Key information to extract from analyzing the response environment should address critical decisions that must be made my decisions-makers, the information needed to make the decisions, and the time the information must be received in order to improve the effectiveness of the response.. Information from the response community delivered in terms of key decisions and timelines coupled with the DBT scenarios will guide the formulation of conceptual architectures to use in generating a set of detection requirements and trade-off curves. An assessment of the current state of technology can be combined with the detection requirements to confirm whether the requirements for a particular component of the architecture are realistic. If the technologies are not compatible or sufficient then the architecture will need to be modified so that the requirements are possible to achieve. Technologies that satisfy the detection requirements can be labeled as favorable technologies to pursue for development.

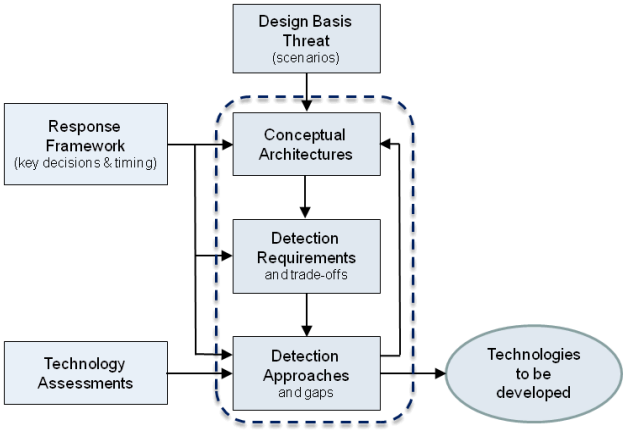

Fig 1. Systems-level view of the major components involved in the process of recommending which detection technologies to develop.

The results and data that subsequent analyses will use to determine requirements across all stages of the detector development process are generated using the following approach. The method begins with the creation of a library of potential attack encompassing a set of credible threat scenarios. Next, a database of response effectiveness is built using information from the response framework and the attack scenario library. Finally, detector architectures are optimized for system performance and the value of the detection system as a function of detector performance parameters is examined.

In the case of an indoor facility, several variables define the attack scenario library including the release location, release size, and airflow state of the building. The exact location of a potential attack cannot be determined *a priori*; therefore, potential release locations are postulated throughout the facility at regularly-spaced intervals. Release sizes are also varied over a credible range that is informed by the threat. The airflow state of the building can be captured by a number of parameters including, but not limited to: the outdoor temperature, HVAC operating condition, and the status of doors inside the facility.

Following generation of the library of potential attack scenarios, a database of the response effectiveness is created by evaluating the mitigating effect of different responses on each attack scenario at various response initiation times. There are a variety of response strategies that could be implemented to mitigate effects of toxic agents, including evacuation, HVAC shutdown, HVAC purge, closing of fire doors, and activation of wet fire suppression systems. The utility of a given response action can vary greatly depending on the properties of the agent and how soon after the release a response is initiated. Figure 2 provides a representative illustration of response effectiveness versus initiation time for a given attack scenario.

Response effectiveness (i.e., the percentage of causalities averted) depends both on the response time and the type of response initiated. "High-consequence" responses, such as full evacuation of a facility, might provide a more effective response to some attacks, but this response is less acceptable when employing a detector system with a high rate of false positives. On the other hand, "low-consequence" responses may involve changing the amount of fresh air in the HVAC system for indoor facilities. While this type of response provides a higher tolerance of false positives, it might result in a less effective response. Moreover, the poor response (red line) in Figure 2 demonstrates that some response actions could even lead to *increases* in casualties (e.g., venting toxic agents to populations located outside a facility or evacuating persons through a cloud of toxic agent). Overall, it is evident that some response actions are more favorable than others, and that there is an increased benefit to initiating the response quickly to maximize the effectiveness. Using FacDAC, a unique response effectiveness curve for each scenario in the library is created and combined into a database. The database of response effectiveness assists in quantifying types of response actions that would be most effective for a given application and over what time frame the response action must be initiated to have a reasonable effect of mitigating the attack.

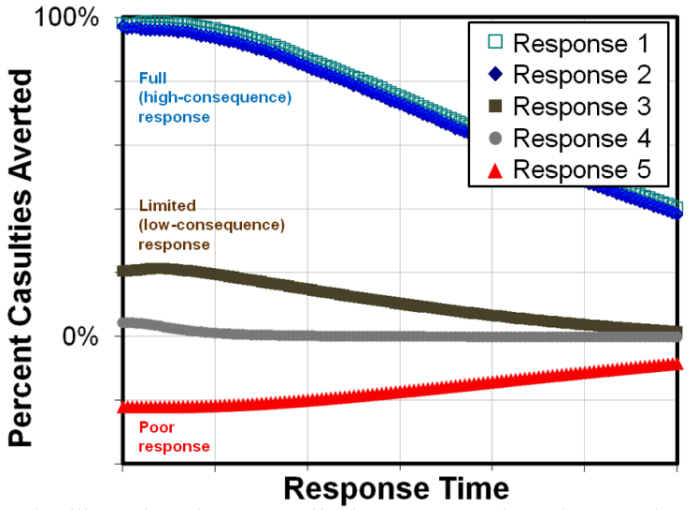

Fig 2. Representative illustration of response effectiveness curves for a given attack scenario. The x-axis indicates the time after the initial release of a chemical or biological agent that a given response

action was taken. The y-axis is the percentage of lives saved relative to performing no response action.

Using the databases of attack scenarios and response effectiveness, conceptual detector architectures are optimized and their effectiveness evaluated. The optimization process for determining detector locations is based on the evaluation of impact-based metrics for various postulated detector locations. The metric discussed in the following example is the number of casualties averted through agent detection and initiation of a response. For a given postulated detector location, the scenarios are individually evaluated to determine if any detector is able to detect the presence of the agent and at what time (post-release) this detection occurs. The number of causalities that result without detection is determined for every scenario in the attack library. For scenarios that were detected, the algorithm determines the number of casualties that still occur even after initiating a given response at the detection time, and the casualty difference between detected scenarios with and without response is recorded. This difference in casualties is the metric used to evaluate the benefit of deploying the detection system at the postulated locations. New detector locations are postulated and the above mentioned process is repeated until the detection architecture that produces the largest mean reduction in casualties is found. At this point the detection architecture is considered optimized.

Using the optimized detection architecture, the value of the system, measured in terms of number of causalities averted, can be evaluated as a function of detector performance parameters, such as limit of detection, number of detectors, and agent properties. The resulting performance curves are used in subsequent analysis, discussed in-depth in the remaining sections of this paper, to determine detector requirements. Having now introduced the approach to evaluate detector performance parameters, discussion in the next section of this paper will demonstrate how the approach can be used to aid in the determination of which technologies to pursue in early-stage detector development process.

4. EARLY-STAGE DETECTOR DEVELOPMENT: TECHNOLOGY INVESTMENT GUIDANCE

Detector system requirements are frequently interrelated and their performance can be traded off against each other to meet specific performance criteria. For example, increasing the sensitivity of the detector could increase the performance of the detection system, but it could also increase the cost of the system. Alternatively, by holding the system costs fixed, one might be able to reach the same level of performance with a greater number of less sensitive but in expensive detectors. Understanding the performance space available for trade-offs between requirements can assist detector designers in potentially excluding from further consideration technologies that demonstrate performance outside the space. The process for determining the trade-space available begins by defining the minimum level of acceptable performance for each requirement of the system. After determining the minimum level of performance additional sections of the trade-space can be examined. For example, there exists a point in the performance parameter spectrum when, for a given application, the benefits of additional improvements in performance will saturate and a diminishing return on investment occurs. Improving the parameter beyond this value will not produce a significant benefit in terms of system performance and represents an upper bound on desired performance.

Although it is useful for detector designers to know the minimum level of performance a system component must possess to be effective, understanding the performance range where the system is expected to give moderate-to-good performance assuming moderate values for other parameters can be even more beneficial. By defining these "reasonable" ranges, boundaries can be placed on the performance space in which trade-offs between specific objectives can be made. Moreover, the method can assist in quickly excluding from further consideration technologies that demonstrate performance below the designated reasonable range. All three important sections of the performance space mentioned above can be succinctly captured and displayed graphically through the use of a construct such as the one provided in Figure 3.

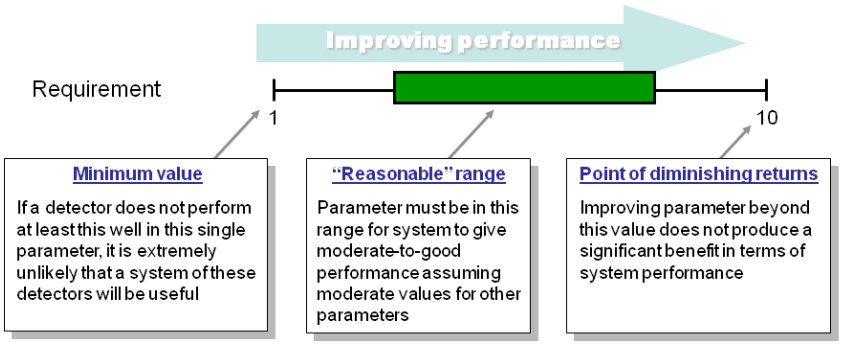

Fig 3. Graphical construct of requirements definition. It is important to note that all parameters within the highlighted (green) range are necessary but not sufficient conditions to meet the performance criteria.

As discussed in Section 3, the value of the detection system as a function of detector performance parameters can be examined by evaluating impact-based metrics of optimized architectures and response effectiveness against a range of credible attack scenarios. An example plot generated by FacDAC is featured in Figure 4a, where the limit of detection (LOD) for a given architecture (e.g., minimum agent concentration that must be detected) is varied and the impact reduction plotted as a function of the number of detectors employed in the detection network. In a previous analysis, not discussed in this paper, it was determined that a reasonable range of three to ten detectors could be expected to be deployed in the system. The point of diminishing returns for the LOD is specified as the initial point of performance saturation for the minimum number of reasonable detectors. In Figure 4a, this occurs at approximately a 90 percent reduction in casualties for the three detector system, yielding a value of 10^{-7}. Furthermore, the minimum value for LOD is specified by examining the minimum acceptable reduction of impact, chosen as ten percent for a ten detector system, yielding a value of 100. Finally, the reasonable range of LOD performance is determined by setting the lower limit of reasonable impact reduction to 50 percent of casualties averted and the upper limit to 90 percent for a given number of detectors in the system, chosen as five in the example. The corresponding graphical requirements construct for the example is presented in Figure 4b.

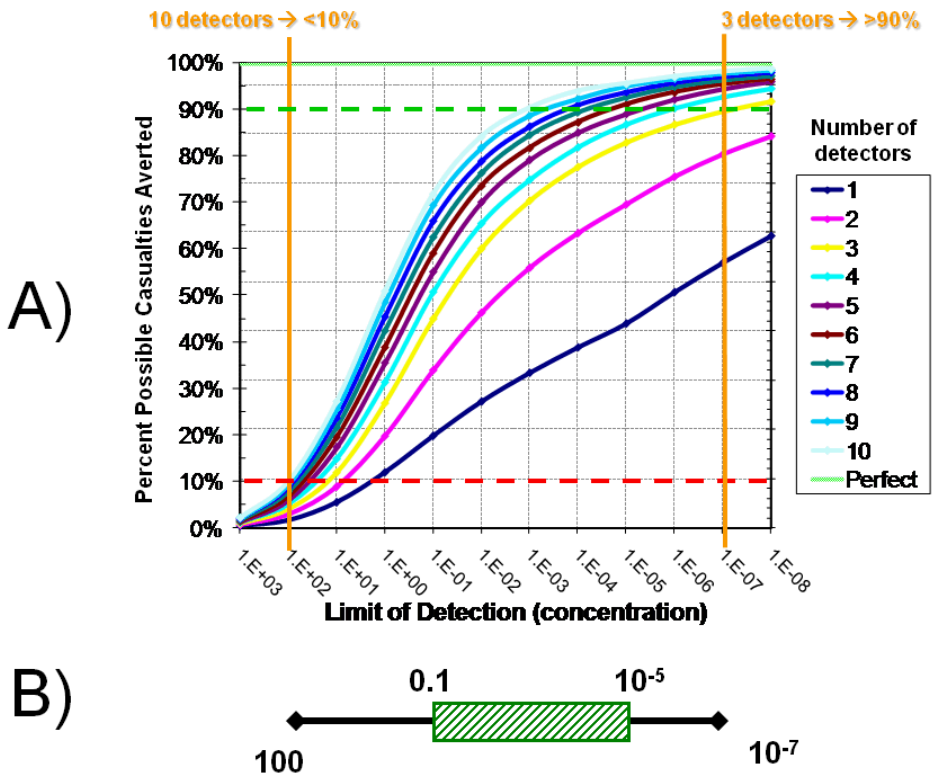

Fig 4. A) Percentage of possible lives saved by detection-triggered evacuation, relative to no action, at various detector limits of detection (i.e., sensitivities) and number of detectors at a given detection time. The orange bars are guides for the eye and highlight the minimum desired reduction in casualties, 10%, and a high reduction in causalities where the benefit of increased sensitivity begins to saturate, 90%. B) Corresponding graphical requirement construct for the LOD.

Creating similar plots for additional performance parameters (e.g., cost) allows for system architecture requirements to be determined. Figure 5 features the results extracted from multiple parameterized performance plots compiled together and displayed as a single system requirement construct. A goal of early-stage detector requirements analysis is to quickly and defensibly rule out the use of unreasonable technologies in the detection system, so that development focus could be placed on more promising technologies. It is important to note at this time that technologies with performance values within the reasonable range may not satisfy the overall system objectives. Existence within the reasonable range is not an indication that the technology meets the desired system performance, only that it is not unreasonable to initially consider detection systems that incorporate this technology. Additional analysis is needed to determine if the individual detector components satisfy the system performance objectives. This type of analysis will be presented in Section 5. Figure 5 also demonstrates how two competing technologies can be evaluated against the system requirements. The first technology, denoted by the black icon, has a more sensitive LOD than the other technology, denoted by the white icon. If the detector system was evaluated solely on the basis of its LOD performance it is clear that both technologies lie within the reasonable range, but the first technology is more sensitive and likely to detect a broader range of attack sizes; particularly at low release amounts. However, evaluation of all the performance criteria shows that the more sensitivity technology has an unreasonable cost for the proposed detection system and should not be considered for this application.

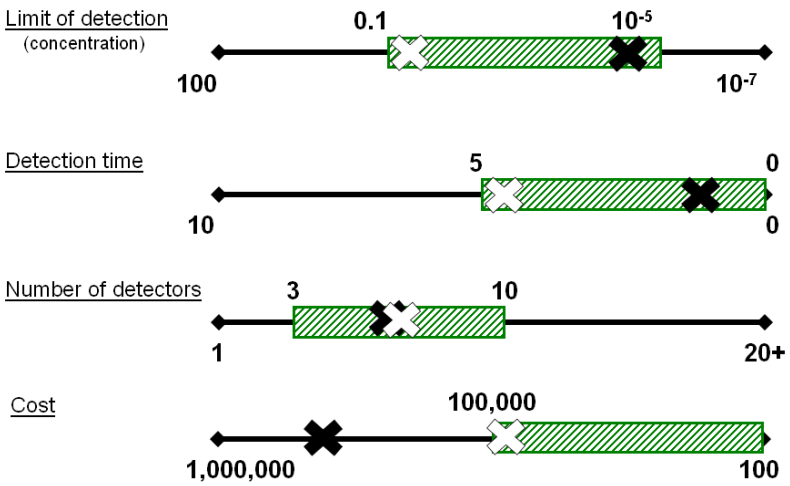

Fig 5. Example chemical detector system requirement ranges. For a given detection architecture, the black and white icons indicate the performance of different technologies under consideration for use in a detection system.

Providing detector developers with requirements shown in Figure 5 will serve as an aid in determining which technologies are most advantageous to pursue. Moreover, the benefit of mapping the performance trade-space of the system allows for the potential discovery of technology or literature gaps. For example, the above illustration suggests that the technology denoted by the black icon could potentially be used in a detection system if the cost was lowered to $100,000. This knowledge would provide defensible guidance to detector designers and program sponsors on where future research and development efforts could be focused.

5. DEVELOPMENT PHASE: EVALUATION OF SPECIFIC SYSTEMS

In the development phase of detector design, after a set of detector technologies has been selected, an end-to-end approach can be used to compare competing technologies and determine whether the detection architectures perform at the required performance level. At this stage of the detector development process, a number of competing technologies could satisfy the minimum and reasonable requirements of the system. What is needed is a rigorous and defensible method to compare the competing technologies against one another and the overall system performance goals. One such method involves taking key detector performance parameters that drive detector performance (e.g., cost, sensitivity, and detection time) and plotting them on an iso-performance chart that displays systems with equivalent performance. An example iso-performance chart, where every point on the plot provides identical system performance, is provided in Figure 6. The lines in the plot were created by extracting data points from performance curves, such as Figure 4a, at a single value of performance (e.g., percent causalities averted). It is important to highlight that each performance curve is generated from the initiation of a specific response at a specific detection time; therefore, many performance curves generated with various detection times are needed to create the iso-performance chart featured below. The iso-performance plot in Figure 6 can be interpreted as follows: a detection system that contains a sensitivity of 5, a detection time of 4, and a cost of $100,000 provides identical system performance to a system that contains a sensitivity of 10, a detection time of 8, and a cost of $25,000. Using an iso-performance plot, the system performance of competing technologies can be quickly and directly compared and each system's ability to satisfy the overall performance objectives can be evaluated.

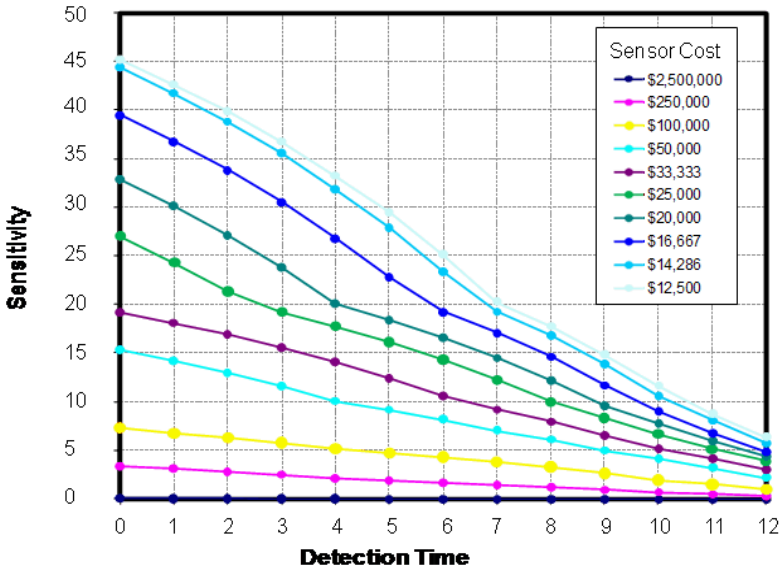

Fig 6. Iso-performance plot depicting a constant performance space for a new detection system. Every point on the plot indicates that the new detection system performs equivalently to a reference detection system. Cost in this figure is a surrogate for number of detectors and would be a combination of acquisition cost and annual operational costs.

6. APPLICATION STAGE: EFFICIENT ALLOCATION OF ASSETS

During the application or deployment stage of detector development, the end-to-end approach can be employed to evaluate the performance of current architectures and calculate optimal detector locations to maximize system coverage. As shown in Figure 2, the most effective response actions are typically "high-consequence"; therefore, having high confidence that a detector alarm indicates a true threat to public health is critically important. One potential method to increase confidence in the system is to optimize detector locations so that a high probability exists that one or more detectors in the network will alarm if an agent is released anywhere within the facility. Figure 7 provides sample outputs of a detector siting optimization generated by the FacDAC toolset. In the figure, the positions for a given number of detectors are optimally sited in order to minimize the largest fraction of population that would be infected (FPI) in an undetected attack. Each point on the figure is an individually optimized detector architecture. It is evident that the benefit of adding additional detectors to the architecture begins to saturate and a point of diminishing returns is reached after deploying six detectors. Marginal benefit tables can be constructed to translate changes in the FPI metric to the number of people protected per additional detector deployed. Informed by these analyses program managers, detector designers, and response planners can better allocate resources to provide the broadest coverage for a given situation.

A)

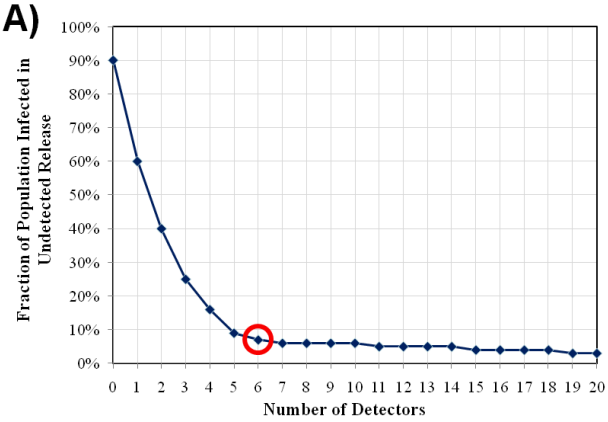

B)

Collector	Fraction Population Infected	Additional Percent Covered	Additional People Protected
1	60.6%	29.4%	14,700
2	40.4%	20.2%	10,100
3	25.4%	15.0%	7,500
4	16.4%	9.0%	4,500
5	9.3%	7.1%	3,550
6	7.2%	2.1%	1,050
7	6.8%	0.4%	200
8	6.3%	0.5%	250
9	6.1%	0.2%	100
10	6.0%	0.1%	50

Fig 7. A) Fraction of population infected for a given number of detectors positioned in optimized locations to provide the greatest amount of coverage in the system. Each individual point represents an optimized detector configuration for a given number of detectors. B) Marginal benefit table showing the number of additional people protected per additional detector deployed. Assumed population of 50,000.

7. CONCLUSION

In conclusion, employing an end-to-end approach to surveillance and detector requirements design can result in robust, defensible, and realistic definitions. The utility of the approach was demonstrated in a number of phases of the detector development process including: providing technology investment guidance during early-stage detector development, assisting in system performance evaluation and down-selection of competing technologies in the development stage, and optimizing asset allocation for chosen detection systems during deployments. Sandia's Facilities Weapons of Mass Destruction Decision Analysis Capability (FacDAC), which captures the end-to-end approach was discussed and applications of its use were illustrated using examples from a number of studies performed for the Department of Homeland Security.

8. REFERENCES

[1] WORLD AT RISK: The Report of the Commission on the Prevention of WMD Proliferation and Terrorism. http://www.preventwmd.gov, 3/9/2009

[2] The BioWatch Program: Detection of Bioterrorism, Congressional Research Service Report No. RL 32152, November 19, 2003;
http://www.dhs.gov/xabout/structure/gc_1205180907841.shtm;
http://www.milnet.com/wh/DoHS/BioWatchFactSheetFINAL.pdf

[3] PROTECT (DHS Program for Response Options and Technology Enhancements for Chemical/Biological Terrorism); http://www.sandia.gov/mission/homeland/chembio/integration/facHarden/protect.html;
http://www.dhs.gov/xres/programs/gc_1217620103978.shtm;
http://www.anl.gov/techtransfer/pdf/Profile_Protect_9-3-04.pdf;

[4] United States Patent Application 20060271211, November 30, 2006