Built on LDRD Laboratory Directed Research and Development

SAND2016-XXXXR
LDRD PROJECT NUMBER: 198564
LDRD PROJECT TITLE: Biosecurity through Public Health System Design
PROJECT TEAM MEMBERS: Walt Beyeler, Patrick Finley, William Arndt, Alex Walser, Michael Mitchell

ABSTRACT:

We applied modeling and simulation to examine the real-world tradeoffs between developing-country public-health improvement and the need to improve the identification, tracking, and security of agents with bio-weapons potential. Traditionally, the international community has applied facility-focused strategies for improving biosecurity and biosafety. This work examines how system-level assessments and improvements can foster biosecurity and biosafety. We modeled medical laboratory resources and capabilities to identify scenarios where biosurveillance goals are transparently aligned with public health needs, and resource are distributed in a way that maximizes their ability to serve patients while minimizing security and safety risks. Our modeling platform simulates key processes involved in healthcare system operation, such as sample collection, transport, and analysis at medical laboratories. The research reported here extends the prior art by provided two key components for comparative performance assessment: a model of patient interaction dynamics, and the capability to perform uncertainty quantification. In addition, we have outlined a process for incorporating quantitative biosecurity and biosafety risk measures. Two test problems were used to exercise these research products examine (a) Systemic effects of technological innovation and (b) Right-sizing of laboratory networks.

INTRODUCTION:

Providing public health services entails the storage, transport, and generation of agents that pose risks to safety and security. In addition, some of the equipment, supplies, and skills developed to advance public health can be diverted to malicious use. There are many approaches to balancing the trade-offs implicit in the need for dual-use materials and technologies in public health services. Educating medical laboratory personnel about risks, and training them in improved practices, is one effective strategy. Inspection of laboratories, resulting in recommendations and implementation funding directed at improving biosafety and biosecurity, is another approach that has a history of success. We have developed a third strategy, which we call systemic biosecurity assessment, that considers the set of institutions and facilities involved in public health as interacting components that serve regional or national demand. This system view lets us pose questions about the overall distribution of resources and capabilities, in particular whether alternative distributions might better serve public health and create a safer and more secure system. This strategy complements existing approaches by insuring that facilities' capabilities

and resources are well-aligned with overall public health needs, while training and inspection of those facilities insures those capabilities and resources create minimal risks.

This report documents a research and development effort used to create essential new elements of our systemic biosecurity assessment platform called Phantom. Phantom is a consolidation of recent work on medical systems logistics modeling, beginning with a study of diagnostic medical sample transport to support control of Ebola in West Africa. Models of facilities and processes central to providing medical care, such as clinics, hospitals, testing laboratories, transport couriers, and patient communities have been developed for the platform, or have been adapted from past related work. The goal of this project was to develop and incorporate models of community demand for medical services, to incorporate uncertainty quantification capabilities which are essential for comparing the performance of alternative system configurations, and to analyze two hypothetical scenarios as a test and demonstration of the new capabilities.

One challenge for a system-level assessment is to select measures for performance that are meaningful for decision-makers. The measures initially defined for the system are focused on laboratory functions, such as the number of samples analyzed and the analysis error rates. In the course of the current work we found that the development of some formal measures for biosecurity and biosafety risk could no longer be deferred. There are many factors regarding facility design, location, and staffing that interact to influence these risks. Thorough exploration of these factors is well outside the scope of this project, however we developed and implemented some simple measures as placeholders and goads for better risk models. This is in keeping with the general philosophy behind Phantom, which is designed as a library of models for basic system processes, which can be selected and configured to address specific problems, and which can be improved over time.

DETAILED DESCRIPTION OF EXPERIMENT/METHOD:

The research described in this report was designed to provide the underpinnings for two enhancements to the Phantom platform's capabilities: representation of community interaction with the healthcare system and introduction of uncertainty quantification techniques to support decision analysis. These capabilities are essential for creating a complete context for comparing the performance of alternative system designs in providing public health. While fostering more efficient public health systems can improve biosafety and biosecurity through more efficient and transparent resource use, the platform does not include explicit metrics for biosafety and biosecurity. In the course of working to close the capability gaps motivating this project, the need for a quantitative assessment of differential security and safety risk became evident. We have therefore set the additional goal of making some preliminary progress in the area. Work done on each of these tasks is summarized below.

Community interaction model

Existing system assessments use a specified demand for sample collection. This demand may vary with location and with time, but it is an exogenous function. Many questions of concern involve assessing the ability of the public health system, including analytical labs, to halt and reverse the progress of diseases that progress over short time periods. Influenza outbreak detection and characterization is an important example. When the disease dynamics are as fast as the medical system response dynamics, sample demand cannot properly be treated as an exogenous process. To do so risks mismeasuring the cost of delay and waste within the medical system. These inefficiencies can feed back to make the underlying disease problem worse, creating an exponential demand that eventually overwhelms response capacity.

One research goal was to develop and implement a community interaction model that can serve as a pattern for later elaborations. We did not opt for an epidemiological model as there are already several suitable implementations available for future incorporation. To define the community/medical system interaction pattern we chose instead to develop a model of an interaction process that had not previously been addressed to our knowledge, and that was directly relevant for some of our application cases.

Successful treatment of some diseases, such as HIV/AIDS and TB, requires administering drugs and monitoring health over an extended period of time. Programs to control these diseases create a long-term and comparatively predictable demand for sample analysis, provided the treatment population is comparatively well defined. Providing analytical results promptly is important because they are needed for monitoring treatment compliance and effectiveness, and for adjusting treatment approaches. Long delays in receiving lab results, or failure to receive them altogether, can lead to ineffective control or to patient drop-out. We reviewed models of customer demand for services to identify a suitable general formalism for this process, implementing it in a format compatible with the existing Phantom framework.

Uncertainty quantification

Model parameter values appropriate for specific scenarios are often highly uncertain. Information about possible demands on the system, operability and capacities of equipment, status of roads, availability of reagents and other key inputs are often difficult to acquire and maintain. Estimates of possible system performance are uncertain as a consequence. Despite this uncertainty, decisions about how resources, equipment, and capabilities will be allocated across the system will be made. Model estimates of performance can contribute to these decisions provided underlying uncertainties are reflected in those estimates. Uncertainty quantification is a set of techniques for reflecting uncertainties about a system's process models and parameter values into probability distributions of possible performance.

Biosecurity and biosafety risk measures

Different configurations of the facilities and resources used to serve public health intuitively present different biosecurity and biosafety risks. For example, distributed laboratory testing may reduce sample transport risk because analytical capabilities are near many points of sample

collection, but increase risk from theft of stored samples because they are held in many sites having variable physical security. Beneficial analysis could therefore be conducted to compare cost differences between transportation security improvements versus facility security upgrades.

Measures of biosecurity risk and biosafety risk can help study these tradeoffs systematically, and to assess the safety and security implications of alternative strategies for expanding public health services. While there are clear benefits to quantifying these risks, unlike the basic system measures we report for laboratory function (number, cost, turnaround time, error rate), developing appropriate biosafety and biosecurity measures is neither straightforward nor uncontroversial.

Developing a single integrated biosecurity and biosafety risk model, appropriate for all contexts and accepted by all analysts, is an unachievable ideal. But a unified approach is not necessary for making significant progress in formalizing comparative risk assessments. The Phantom platform is designed to provide a library of models for the basic processes in the system (e.g. sample collection, transport, and analysis) which can be selected and configured as the analytical context requires. This same architecture will support the definition of multiple risk models, which analysts can select from and tailor based on their knowledge and experience. In addition, the uncertainty quantification capability added to Phantom as part of this research project supports application and assessment of alternative risk models. Parameters which may not be well known, such as the probability that a lab technician might become an inside agent, can be assigned a wide range of possible values as part of an overall uncertainty analysis. This approach allows analysts to assess whether different configurations pose significantly different risks, given their uncertainties about the relevant factor values, and whether those differences persist across the various models they consider plausible.

Given the complexity of the problem and the limitations on time and resources of the current project, our goal is not to propose a dispositive model of biosecurity and biosafety risks, but to develop an initial simple model as a way of including risk quantification as a new class of formal metrics available for describing the system. Future research can build on this model, and add alternatives for analysts' use, once this ground has been broken. Our approach was to survey current work in the area, and to define the basic computational infrastructure required for incorporating safety and security metrics into the Phantom framework.

RESULTS:

Research performed during this project led to model enhancements in the critical areas of community/medical system interaction and uncertainty quantification. An initial risk quantification framework was also added to support introduction of biosafety and biosecurity models. These new capabilities were then exercised in two sets of stylized analyses involving alternative laboratory configurations in Cambodia. Results are summarized below.

Community interaction model

We use a Markov model to describe patients' state with respect to a disease management program. The model generalizes over diseases, as is designed to apply to a range of conditions that require periodic monitoring of a patient population, with the population growing or declining over time. Figure 1 shows the population groups considered and the transition rules and processes that connect them.

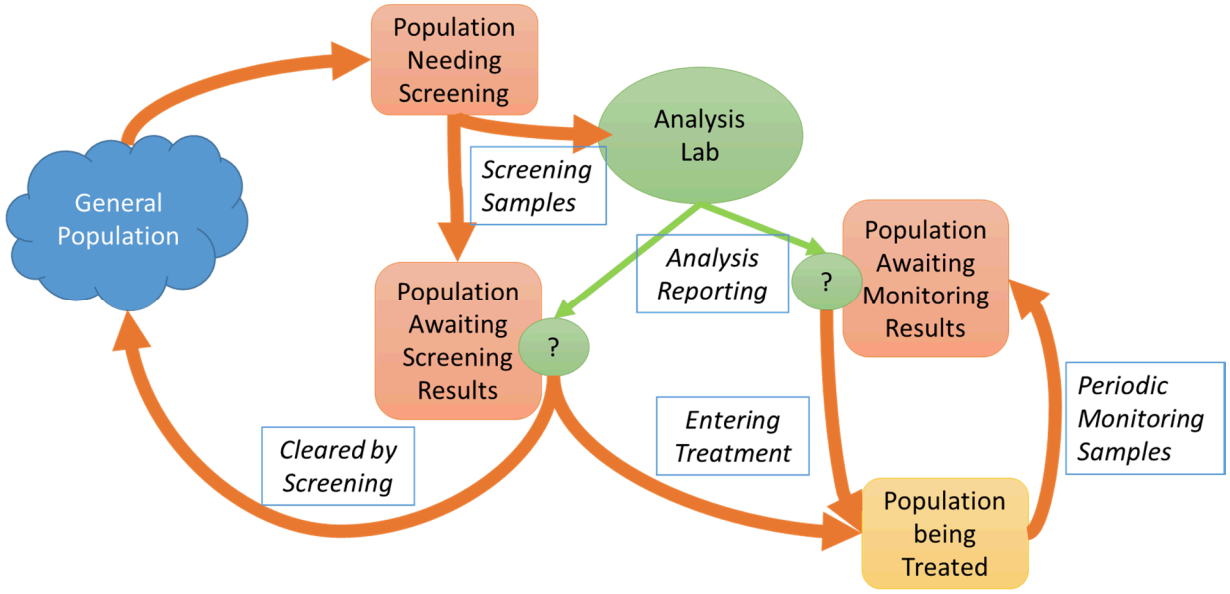

Figure 1 - Interacting Patient Flows and Information Flows in the Community Response Model

Prospective patients emerge from the general population with a specified frequency. A screening sample is required to determine whether a prospective patient should receive treatment. If the screening result is positive the prospective patient is added to the population being treated. Each member of the treated population is periodically monitored to assess the effectiveness of their treatment. Their specific test results may lead to changes in their course of treatment. Here we assume that they remain under some treatment regimen, and so are liable for future monitoring regardless of the test outcome. (A straightforward elaboration of the model would allow for patients to complete treatment based on results with some probability). The aspect of the test we focus on here is the timeliness of the result rather than its specific value. The longer results take to be reported back to the requesting physician, the more likely it is that the patient's records will have been misplaced, the patient will have returned home or moved, or will have become "lost" to the treatment program for some other reason.

We model this process as a random loss rate with a characteristic time constant that operates on the population of patients waiting for sample results. Different time constants can be specified for screening and monitoring tests. If the information comes back quickly, considering transport times, queueing time in the lab, actual time required to conduct the test, and the efficiency of the notification system, the patient loss rate is small. If aggregate delays become as large as the loss time constant, a significant fraction of patients needing treatment will not receive it. The integrated patient loss rate therefore becomes another potential measure of system function.

Uncertainty quantification

The ability to repeatedly execute a simulation model using alternative values of key parameters is essential for any sample-based uncertainty quantification. The initial use case for Phantom was for interactive execution through a web interface. This mode allows decision-makers and analysts to experiment with scenario parameters and directly observe results. The Phantom architecture cleanly separates the user interface service from the simulation model service. It was designed in this way so that the simulation engine would also be available for the repeated executions required for uncertainty quantification.

Research accomplished during this project integrated existing capabilities for massively parallel execution of simulation models with the Phantom simulation engine interface. This involved three software development tasks. First, the model configuration algorithm was modified to allow symbolic specification of a broader range of parameter values. Second, a specific driver code was developed to translate inputs from the existing generic model driver into symbol/values pairs to be used by the configuration engine. Third, a file-based output capability was added to facilitate parallel execution on HPC resources. This prevents the database connection contentions that would otherwise limit execution speed if model outputs were directly committed to the database, as is done for single interactive simulations.

Biosecurity and biosafety risk measures

Sandia National Laboratory's Biorisk Assessment Model (BioRAM) has been identified as a suitable system to conduct biosecurity and biosafety analyses. Designed for use by biorisk officers at laboratories, BioRAM provides visualization of relative risks, and also helps to identify risk mitigation measures. The assessment process is broken into components:

- Evaluate the biological agents that exist at the facility.
- Evaluate the facility processes and procedures.
- Evaluate the in place biorisk mitigation measures.
- Evaluate the potential adversaries of the facility.

Within each component are several criteria and sub-criteria that are scored independently. These scores are weighted and then rolled up to provide the overall consequence and likelihood score. This method is based on a Multi Criteria Decision Analysis (MCDA) scheme, quantifying the various aspects of biorisk using qualitative definitions. The final results show the relative risk of agents at the given facility, and give program management a mechanism to determine risks that are unacceptable. This scheme can aid program management in allocating recourses to mitigate facility biorisks; or to assess current biorisk program management effectiveness.

BioRAM analyses can therefore provide the requisite information for the Phantom model, giving the user a more detailed understanding of how to improve biosecurity at a national and/or regional level. For example, it may initially appear that security of a given region could be increased by routing samples to more proximal and/or lesser used facilities, reducing travel times

and distances. However, should these facilities lack rudimentary security infrastructure, samples would end up being stored at an unsecure location, reducing overall security. The BioRAM model could help distinguish security lapses at the facility and area level, granting the Phantom user a more complete picture of the most secure means on which to manage a national or regional biorisk management program.

Test problems

Two test problems were defined to exercise the population response models and uncertainty quantification approaches developed above in the context of a realistic system analysis. Both cases involve comparing system-level performance under scenarios corresponding to alternative decisions available to medical system funders or operators. A brief description of each test problem is followed by a discussion of comparative system performance under those alternatives.

Both problems involve hypothetical modifications to the collection and analysis of medical samples in Cambodia. Past project engagement with the Cambodian Ministry of Heath has provided an understanding of baseline conditions sufficiently detailed to enable a realistic analysis. However, to protect information provided to us through that work, data used in the analyses reported here has been de-identified and altered so that it no longer accurately describes existing facilities. These results should therefore not be used as a basis for decision-making.

To support treatment and control of HIV AIDS, tens of thousands of viral load tests are performed annually. Current practice involves transporting whole blood samples for analysis using Abbott m2000 instruments in Phnom Penh. This practice entails several logistical complications, include coordinating sample collection, coordinating transport, and preserving cold chain. Near-POC alternatives, using GXS machines at or near collection facilities, are under consideration. This design would speed turnaround (and therefore simplify follow-up) and remove some logistical problems. Resupply and maintenance logistics and cost may be more complicated, and test quality may suffer. We compare the performance of the two strategies using a range of parameter values to characterize uncertainties about test performance.

This scenario does not directly address a significant biosecurity concern. However, it demonstrates the analytical tools needed to measure tradeoffs between centralized and distributed networks for analyzing medical tests. This general framework is directly transferrable to applications in which differential biosecurity risks among analysis techniques is one of the performance metrics.

We first consider a baseline scenario in which three Abbott m2000 machines are used to analyze whole blood samples. Two are located at the NCHADS laboratory near Phnom Penh, and one at the provincial hospital lab in Siem Reap. In this configuration samples taken at the 51 treatment centers are transported to the nearer of the two labs. Parameters describing the scenario are given in Table 1. Several of these parameters are uncertain, and were therefore assigned probability distributions to cover a range of possible values.

Table 1 – Value ranges for parameters sampled in the test comparison of alternative HIV laboratory networks	
Parameter	Range
Time constant for patients abandoning the program	0 – 4 weeks
Critical fractional utilization of testing equipment for onset of congestion errors	0.5 - 1
Incident rate of new potential cases	0 – 0.5 per day
Abbott m2000 characteristics	
Nominal error rate	0
Maximum false positive error rate	0 – 5%
Maximum false negative error rate	0 – 5%
Xpert MTB/RIF characteristics	
Nominal error rate	8%
Maximum false positive error rate	8 – 20%
Maximum false negative error rate	8 – 20%

For each configuration, 1000 simulations were made, each using a vector sampled from the parameter space described in Table 1. Figure 2 shows cumulative probability distributions of two summary statistics useful for comparing the performance of the two configurations. Cost per sample is especially salient in low-resource countries. The number of samples collected depends on the number of patients under treatment, the arrival and removal rate from this population, and the patient monitoring period. The chance that a patient's sample would be misclassified (either a false positive or a false negative) is another important figure of merit. Analysis error is modeled as a function of baseline accuracy (assumed to be perfect in the case of the m2000) and of system load.

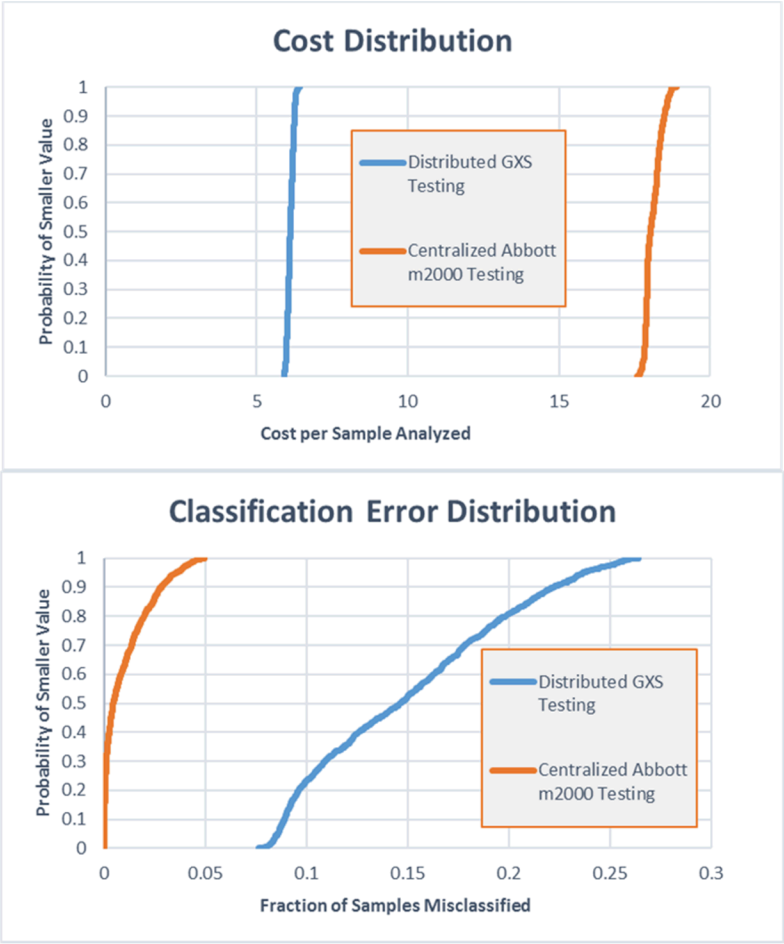

Figure 2 - Distributions of Analysis Costs and Error Rates for Two HIV Monitory Networks

Clinical tuberculosis samples are currently cultured and analyzed at three laboratories in Cambodia. Consolidating analyses at two locations would potentially reduce risk and lead to more efficient use of resources and laboratory personnel. Surge capacity and system reliability would be sacrificed, and turnaround time and transport costs would be likely to increase. We simulated a consolidation scenario in which TB testing was stopped at one of the three locations.

Figure 3 shows compares error rates and analysis turn-around times for the two scenarios. Turnaround times are systematically longer in the two-lab system vis-à-vis the three-lab system. Sample time disributions for each scenario are very narrowly distributed because this analysis did not include stochasticity in sample generation: and so both scenarios are presented with the same stream of samples in each realization. Distributions of error rates are substantially identical in the two scenarios. Under the simulation conditions, lab consolidation would still lead to acceptable performance.

Figure 3 - Distributions of Analysis Turn-around time and Error Rates for two TB Testing Scenarios

DISCUSSION:

The community interaction model developed as part of this project captures a simple form of feedback between the performance of the public health system and the community. While helpful for measure some of the consequences of prompt sample turn-around, as illustrated in the HIV scenario analysis, its larger value was in allowing us to introduce a new feedback mechanism into the modeled dynamics. Representing the coupling between community state and the functioning of the medical system is essential for assessing response strategies that can effectively respond to public health conditions with rapid growth potential. ILI outbreaks are a salient example. Elaboration of the current model architecture to analyze monitoring strategies for emerging infectious diseases has been greatly advanced by the present work.

Incorporation of uncertainty quantification into our analysis process was essential for creating a reliable decision-support tool. The initial, interactive, mode of exploring alternative system designs served many objectives well, including stakeholder engagement, review and correction of system information, and development of aggregate measures of system function. We expect this mode of interaction to be of continued importance in these areas, as well as in helping to elicit scenario definitions and possible control strategies from system operators, and for visualizing key simulation results. However, the public health systems of resource-poor countries are very difficult to characterize, and their future condition and performance is subject to a wide range of uncontrollable factors. In order to responsibly evaluate the possible impacts of alternative system configurations it is essential to reflect both epistemic uncertainty and stochastic processes in that evaluation. The incorporation of uncertainty quantification techniques as part of this research has completed an essential step in solidifying our methodology to systemic risk control.

The first necessary step was to create a new driver application for the simulation engine that manages the configuration and execution of the model, and organizes its output. A clean separation between the simulation engine and the part of the application server that manages user interaction was incorporated in the design from the outset in anticipation of supporting access by automated processes in addition to human users. A suitable driver application was designed and developed, as part of this research, to exploit this potential. Extensive prior work in complex systems simulation modeling has produced a library of resources for efficiently exploring model parameter space and for analyzing the information-rich response surfaces and time series produced by this exploration. The Phantom driver application was therefore designed to provide an interface between existing resources for uncertainty quantification and the specific requirements of the Phantom model.

Fulfilling this goal was complicated by the requirement that the driver application be applicable to all Phantom model scenarios, and not be tailored for a specific system of class of analyses. The number and interpretation of parameters used to define a scenario can be arbitrarily different from those used in another scenario. A good solution is for the process controlling the uncertainty quantification process to define a set of symbol/value pairs to the model configuration engine. This frees the interface between them from any dependency on model parameterization. However, this solution imposes a requirement on the model configuration engine to interpret specification via symbol/value pairs. Creating this capability constituted a second major task.

A final task concerned management of simulation output. In its interactive mode, the application server commits all model output to a central database management system (DMBS). This approach is well-suited for providing distributed access to a common repository of results from single simulations. Uncertainty quantification simulations have very different requirements. The large number of simulations, typically produced concurrently on massively-parallel machines, generate large bursts of data that require careful management of database connections and configuration to avoid swamping the DBMS. An effective solution is to divert the output of

individual simulations to their own flat archive on the MPP file system for subsequent importation into the database. This required creation of additional archiving objects.

Explicit quantification of biosecurity and biosafety risk was not an initial objective of this research, however it was seen to be an additional important step towards comprehensively enabling risk management through system redesign. Methodologically, we identified existing tools for analyzing the factors contributing to risk and soliciting and organizing relevant site information. In the Phantom system simulations, processes that contribute to biosafety and biosecurity risk (such as sample collection and packaging, sample transport, preparation and analysis, storage, acquisition of equipment and reagents) are explicitly represented. While the risk model(s) that connects these basic activities with overall risk measures is yet to be developed, we were able, as part of this project, to add appropriate monitoring and accounting objects to collect information about these activities. These basic elements will provide basic information needed for subsequent risk modeling, whatever form that may take,

ANTICIPATED IMPACT:

Sandia is developing a unique capability for bringing system analysis methods and insights to bear on the problem of fostering biosecurity and biosafety. In additional to general analytical resources drawn from diverse engineering applications, this capability draws on the expertise and insights developed through years of engagement with diverse health programs, including detailed understanding of their organization, assets, and capabilities. This ability to integrate relevant facility-level operational information into an integrated system view, and to do so in a simulation environment that enables alternative system designs to be posited and assessed, has had significant success in initial applications.

These applications were made in response to specific needs, and so the support data processing and analytical infrastructure was necessarily ad hoc. More recently, the many artifacts produced in the course of developing this capability have been integrated into a unified simulation platform (Phantom). This integration has also allowed us to incorporate more rigorous software engineering, allow web deployment in a variety of configurations, and provide for data security in a multi-user environment. Development of the Phantom platform has greatly reduced to effort required to develop and deploy models of specific systems, to make those models available to a wide range of users, and to tailor the access of those users to their specific responsibilities.

The value of assembling a system-level view of public health infrastructure, and of simulating its performance under a variety of historical and hypothetical conditions, has been widely appreciated and endorsed by public health system experts in the US and abroad. This unique capability is recognized as enabling great improvements in the way analytical resources are deployed in support of public health, and of understanding the various costs imposed by current deployment practices, including mortgages on public health budgets from equipment maintenance, training, and supply. Biosecurity and biosafety risk are widely understood to be influenced by how and whether these costs are met.

The work performed under this project has filled in essential technical pieces required for us to strengthen our engagement with CDC and other public health agencies. The interaction between the community and the public health system is an essential determinant of the effectiveness of outbreak detection and control. Assessing the relative performance of alternative surveillance and monitoring systems requires the ability to model the dynamics of community interaction with the medical system. These dynamics determine, for example, whether the system can detect and control an infectious disease more rapidly than the disease can spread through the community. Neglecting these dynamics prevents assessment of the likelihood that a given system will be overwhelmed, and therefore precludes risk management.

Our initial community interaction model has allowed us to develop and test the necessary feedbacks. The dynamics we considered do not allow exponential growth of demand, but they do address a phenomenon of concern to program managers considering alternative testing strategies for HIV. Completion of this model introduces a qualitatively new capability, and directly enables engagement on with field researchers. Including an infectious disease process model is an important next step along this direction. Such a capability would allow assessment of EID monitoring system designs, an area of active interest to current and prospective customers.

This project has allowed us to create a core capability to perform uncertainty quantification (UQ) for any model scenario developed in Phantom. This core capability allows us to compare the performance of alternative system designs while formally incorporating limitations on available information and any stochastic processes influencing the system. Creating this ability was a necessary precursor to continued collaboration in evaluating comparative performance of fielded systems. While the basic information management and simulation control components are in place, there is currently no convenient method for defining parameter distributions and sampling designs, nor for portraying analysis results in a convenient web interface. These elaborations may be pursued, as guided by requirements of collaborators and other stakeholders.

We have laid the groundwork for incorporating quantitative risk measures for biosecurity and biosafety. This goal was not anticipated at the outset of the project, but its importance warranted us taking some preliminary steps toward its realization. First, the essential and complex job of analyzing these risks, in the sense of identifying contributing factors and describing their interactions, has been addressed by an existing tool. Some of these factors connect with processes or conditions simulated in the Phantom model. This connection prospectively allows the differential risks associated with alternative system configurations to be derived and made available to the user. We have taken the initial step of creating model objects to monitor the simulation processes that contribute to risk. Subsequent steps would involve integrating the concepts and potentially technologies from the BioRAM tool to develop appropriate quantitative risk measures.

CONCLUSION:

This project has allowed us to do the research necessary to make essential advances in our development of an overall methodology for systemic risk assessment and control. Lack of models of the interaction of the public health system with the population of patients prevented our properly posing the overall control problem, blocking our ability to properly consider scenarios in which prompt reaction is key. New model elements developed here have closed this important hole. Inability to do comprehensive uncertainty quantification prevented our producing technically sound comparisons to decision-makers. Work done here has connected the Phantom simulation platform to a wealth of existing UQ resources at Sandia. While the more efficient design of public health systems, in the sense of insuring that resources, equipment, and skills are transparently aligned with public health needs, intuitively improves biosecurity and biosafety (for example by minimizing idle or underused assets, providing a workflow to labs that develops expertise) the extent of this improvement could not be quantified in a way that informed decisions about alternative system designs. This project has allowed us to take important initial steps in this direction.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.