

Results of a Simulation of Response to a Bioterrorism Crisis

**James E. Coolahan, Ph.D., Michael T. Kane,
Katherine L. Morse, Ph.D., John F. Schloman**

**Johns Hopkins University Applied Physics Laboratory
Laurel, MD**

**James.Coolahan@jhuapl.edu, Michael.Kane@jhuapl.edu,
Katherine.Morse@jhuapl.edu, John.Schloman@jhuapl.edu**

**Roy P. Koomullil, Ph.D.,
Alan M. Shih, Ph.D.**

**University of Alabama at Birmingham
Birmingham, AL**

**rkoomul@uab.edu
ashih@uab.edu**

**Evangelos I. Kaisar, Ph.D.
Florida Atlantic University
Boca Raton, FL
ekaisar@fau.edu**

**Kenneth K. Walsh, Ph.D.
Florida A&M University
Tallahassee, FL
walsh@eng.fsu.edu**

ABSTRACT

The National Center for the Study of Preparedness and Catastrophic Event Response is a consortium led by The Johns Hopkins University (JHU) for the U.S. Department of Homeland Security. One of the projects, the Modeling and Simulation Integration Framework, is focused on preparing for catastrophic events. This project is led by the JHU Applied Physics Laboratory, in collaboration with the University of Alabama at Birmingham, Florida Atlantic University, Florida A&M University, and the Brookings Institution. The first prototype, simulating the response to an urban chemical disaster, was reported upon in a technical paper at IITSEC 2008. The second prototype simulates the longer term response to a bioterrorism crisis triggered by the surreptitious release of smallpox in a large indoor arena during a concert.

This paper describes the design of the Bioterrorism Crisis Management simulation, and the systems engineering considerations that led to structuring the simulation into a set of components executing in a much-faster-than-real-time High Level Architecture federation, and a set of components that are executed prior to federation execution. The components discussed include simulations of transport of the smallpox agent within the arena's ductwork; 3D airborne transport of the agent throughout the arena during the concert, resulting in infection of many attendees; spread of the disease over several weeks in a major metropolitan area; traffic flow around hospitals due to people seeking treatment and vaccination; risk communication effects; and surge capacity for hospital treatment and vaccination.

The paper provides the results of the initial execution of the simulation that show the spread of the disease over time and the results of the medical response, based on the initial conditions and assumptions. Finally, potential application of the simulation to mission rehearsal by emergency response decision-makers is discussed.

ABOUT THE AUTHORS

James E. Coolahan is a Program Manager for Modeling and Simulation, and a National Security Studies Fellow in the National Security Analysis Department at the Johns Hopkins University Applied Physics Laboratory (JHU/APL), where he also served as the Assistant to the Director for Modeling and Simulation (M&S) from 1996 to 2001. In his 37 years at JHU/APL, his technical activities have included M&S, test and evaluation, and the development of oceanographic data acquisition systems. Dr. Coolahan served as a member of the National Research Council (NRC) Committee on Modeling and Simulation Enhancements for 21st Century Manufacturing and Acquisition from 2000 to 2002. He is currently the Principal Investigator for the Modeling and Simulation Integration Framework project being performed as part of the National Center for the Study of Preparedness and Catastrophic Event Response (PACER) for the Department of Homeland Security (DHS). Dr. Coolahan holds B.S. and M.S. degrees in aerospace engineering from the University of Notre Dame and the Catholic University of

America, respectively, and M.S. and Ph.D. degrees in computer science from JHU and the University of Maryland, respectively.

Michael T. Kane is on the Senior Staff of the Modeling and Simulation Group in the National Security Analysis Department at JHU/APL. Mr. Kane has supervised a Software Engineering Section in the Department and has served as the technical lead for projects involving DoD High Level Architecture (HLA) simulations, C4ISR architecture modeling, management information systems and knowledge management. He served as chairman of the Design Steering Committee for the Department of Transportation nation-wide vehicle database management system and as manager of the Advanced Computing and Visualization Laboratory. Mr. Kane holds a B.S. degree in industrial engineering from the University of Pittsburgh, an M.S. degree in operations research from the George Washington University, and an M.S. degree in computer science from JHU.

Katherine L. Morse is a member of the Senior Professional Staff at JHU/APL. She received her B.S. in mathematics (1982), B.A. in Russian (1983), M.S. in computer science (1986) from the University of Arizona, and M.S. (1995) and Ph.D. (2000) in information & computer science from the University of California, Irvine. Dr. Morse has worked in the computer industry for over 25 years, specializing in the areas of simulation, computer security, compilers, operating systems, neural networks, speech recognition, image processing, and engineering process development. Her Ph.D. dissertation is on dynamic multicast grouping for data distribution management, a field in which she is a foremost expert. Dr. Morse was the 2007 winner of the IEEE Hans Karlsson Award.

John F. Schloman is a Software Engineer at JHU/APL. In 2003 he completed his Master's in Computer Science at Michigan State University. In 2001 he received a degree Cum Laude in Systems Analysis from Miami University in Oxford, OH.

Roy P. Koomullil is an Associate Professor of Mechanical Engineering at the University of Alabama at Birmingham (UAB). Dr. Koomullil received his doctoral degree in Aerospace Engineering from Mississippi State University (MSU) in 1997. Before joining UAB, he worked as an Assistant Research Professor at the NSF Engineering Research Center (ERC) at MSU. Dr. Koomullil's area of expertise includes compressible and incompressible flows, computational fluid dynamics, finite volume upwind schemes on generalized grids, unsteady flow simulation, six-degree-of-freedom (6-DOF) rigid body simulations, overset meshes, and high performance computing. He is a Senior Member of the American Institute of Aeronautics and Astronautics (AIAA) and a Member of the American Society of Mechanical Engineers (ASME).

Alan M. Shih is a Research Professor and the director of the Enabling Technology Laboratory (ETLab) in the Department of Mechanical Engineering, UAB. Dr. Shih received his Ph.D. degree in Aerospace Engineering from MSU in 1994. His areas of interest include numerical geometry, mesh generation, computational graphics, scientific visualization and virtual environments. He is an Associate Fellow of the AIAA and a Member of the ASME and IEEE.

Evangelos I. Kaisar is an Assistant Professor at the Florida Atlantic University. Prior to joining the University he was a Projects Manager of the Traffic Engineering division in the Maryland Transportation Authority (MdTA). In MdTA, his technical activities included traffic and congestion pricing studies, simulation and modeling, and safety studies. Dr. Kaisar served as a member of the Transportation Research Board (TRB), Committee for Intermodal Terminals, Port Operations, Network Equilibrium and Safety and Security. He is currently the Principal Investigator for the following projects: a) Simulation modeling project being performed under JHU/APL for DHS, b) Intermodal Safety and Security project being performed for the Federal Transit Administration, and c) Congestion impact study for the Port of Miami. Dr. Kaisar holds a B.S. degree in civil engineering from the University of Maryland and National Technological University of Athens, Greece, and M.S. and Ph.D. degrees in Civil Engineering, minor in Logistics, from the University of Maryland.

Kenneth K. Walsh is currently serving as a Research Associate in the Center for Thermal Energy Management Progress (CTEMP) at Florida A&M University. Dr. Walsh's research interests center around structural vibration control with particular focus in the areas of semi-active control, innovative materials for improved structural performance, and vibration absorption and isolation of mechanical systems. Dr. Walsh holds a B.S. degree in Civil Engineering from Florida State University and M.S. and Ph.D. degrees in Civil Engineering from Florida A&M.

Results of a Federated Simulation of Urban Chemical Disaster Response

**James E. Coolahan, Ph.D., Michael T. Kane,
Katherine L. Morse, Ph.D., John F. Schloman**

**Johns Hopkins University Applied Physics Laboratory
Laurel, MD**

**James.Coolahan@jhuapl.edu, Michael.Kane@jhuapl.edu,
Katherine.Morse@jhuapl.edu, John.Schloman@jhuapl.edu**

**Roy P. Koomullil, Ph.D.,
Alan M. Shih, Ph.D.**

**University of Alabama at Birmingham
Birmingham, AL**

**rkoomul@uab.edu
ashih@uab.edu**

**Evangelos I. Kaisar, Ph.D.
Florida Atlantic University
Boca Raton, FL
ekaisar@fau.edu**

**Kenneth K. Walsh, Ph.D.
Florida A&M University
Tallahassee, FL
walsh@eng.fsu.edu**

BACKGROUND

DHS PACER Center of Excellence Background

In late 2005, the U.S. Department of Homeland Security (DHS) selected a multi-institution consortium led by the Johns Hopkins University (JHU) to form a National Center for the Study of Preparedness and Catastrophic Event Response (PACER). The Center is performing research into preparation for high-consequence events, and its research is addressing the technical, systemic, behavioral, and organizational challenges such events pose.

M&S Integration Framework Project Background

As part of its research program, the PACER Center has established several integrated cross-cutting projects, one of which is the Modeling and Simulation (M&S) Integration Framework, focused on preparing for the response to catastrophic events. This three-year project, started in the fall of 2006, is led by the Applied Physics Laboratory of JHU (JHU/APL), and involves researchers from the University of Alabama at Birmingham (UAB), Florida Atlantic University (FAU), Florida A&M University (FAMU), and the Brookings Institution. The M&S framework is intended to provide a composable set of simulations that can be used as an aid to decision-makers in training/rehearsal for responses to catastrophic events.

As the first prototype to demonstrate the initial operating capability for the M&S Integration Framework, the four-institution team of JHU/APL, UAB, FAU, and FAMU undertook the development of

an Urban Chemical Disaster (UCD) simulation, which was designed to simulate the emergency response to an urban chemical disaster – the release of chlorine from two explosively ruptured railcars near a downtown area (Coolahan et al, 2008).

THE BIOTERRORISM CRISIS MANAGEMENT SIMULATION – OBJECTIVES AND SCENARIO

Objectives

To demonstrate the framework on a long-duration scenario, a Bioterrorism Crisis Management (BCM) simulation was selected, the objectives of which are:

1. To demonstrate a second prototype simulation for the PACER M&S Integration Framework.
2. To simulate the release and airborne transport of a dangerous biological agent in an indoor arena during a concert attended by over 10,000 people.
3. To simulate the spread of the resulting communicable disease over multiple weeks in a large urban area and its nearby suburbs.
4. To simulate a realistic flow of the population from home to work and school during the outbreak.
5. To represent the command and control, and risk communications, that might be undertaken once the outbreak begins to be recognized.
6. To simulate hospital surge, both for treatment, and for vaccination of portions of the population.
7. To simulate the traffic flow that could result in the vicinity of treatment and vaccination centers.
8. To ensure that the simulation federation can execute much faster than real time on a small set of personal computers.

This research was supported by the U.S. Department of Homeland Security (Grant Number N00014-D6-1-0991) through a grant awarded to the Center for Study of Preparedness and Catastrophic Event Response (PACER) at the Johns Hopkins University. Any opinions, findings, conclusions and recommendations expressed in this publication are those of the author and do not represent the policy or position of the Department of Homeland Security.

Scenario

The scenario begins with a popular music concert held in a downtown arena in the City of Baltimore, which was chosen because of the familiarity the team gained with the city during the UCD simulation. Baltimore has an arena often used as a concert venue, but it was built around 1960 (and is planned to be replaced shortly), and the building plans available to the team were sparse and manually drawn. So, a different arena with more available and modern plans was identified and “geo-relocated” to Baltimore for the scenario. Configured for a concert, its capacity is over 11,000.

Near the beginning of the concert, smallpox is surreptitiously released from a container in aerosolized form in part of the air handling system. Smallpox was chosen because of its potential lethality and because symptoms will not appear for a week, making its recognition and the evaluation of the source of the outbreak difficult. Over the next two hours, the crowd is exposed to various concentrations of the biological agent, depending upon their seat location and the movement of air within the arena.

The concert-goers then return to their homes in the local metropolitan area and resume normal activities. To keep the simulation scale manageable, it was assumed that all concert-goers reside in an area roughly bounded by the Baltimore Beltway, which encircles the city. After about a week, non-specific symptoms, followed by specific symptoms, begin to appear in some of the concert-goers. The disease begins to spread based on home-work-school contact patterns. Within a day or so of the recognition of smallpox-specific symptoms in several persons, local medical response begins. A command and control

system is activated, the nature of the outbreak is assessed, risk communications are initiated, and instructions are given to the population regarding seeking treatment and getting vaccinations. The population reacts, based on the risk communications, media reports, and their own beliefs. Traffic to treatment centers builds in response to the population’s reactions. The spread of the disease and the local responses continue over several weeks.

BCM SIMULATION DESIGN OVERVIEW

Figure 1 shows a simplified block diagram of the BCM simulation. Both the simulation of the airflow within the arena and the transport of the smallpox in that airflow require computations that currently cannot be performed in real time. Therefore, both the simulation of the release through the arena ductwork and the simulation of the airborne transport within the arena must be performed in advance. Also, the generation of the local population, its age demographics, home and work/school locations, and the subset attending the concert, can be performed in advance.

Once these computations are performed, the faster-than-real-time BCM simulation federation can be executed, which consists of four principal simulation components:

- the spread of the disease through the local population and the population’s decision-making;
- the flow of traffic, both in daily home-work-school patterns and to treatment centers;
- the command and control of response actions, including risk communications;
- the surge of persons to be treated at hospitals and vaccination centers.

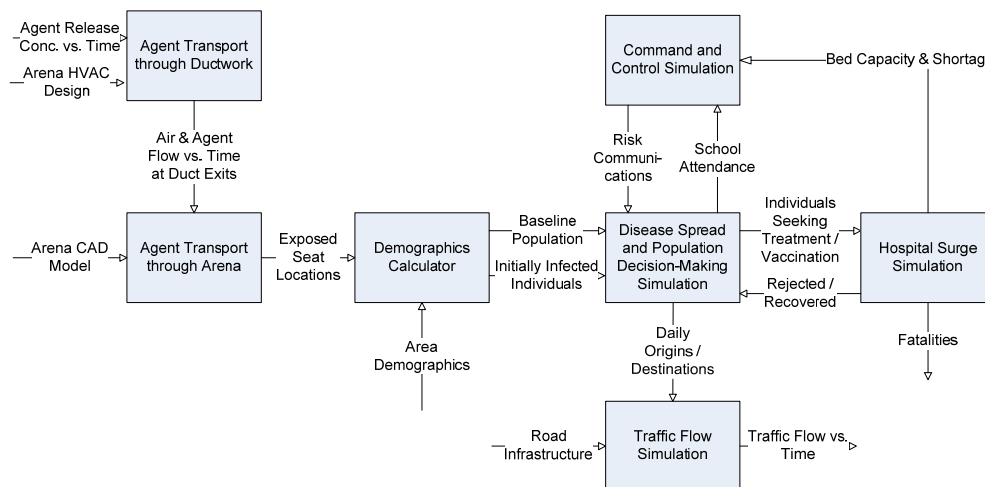


Figure 1. Simplified Bioterrorism Crisis Management Simulation Block Diagram

PRE-EXECUTED SIMULATION COMPONENTS AND RESULTS

Agent Transport in Arena Ductwork

The simulation of the transport of the smallpox from its release through the arena ductwork was performed by researchers at FAMU using the CONTAM simulation. It is a multi-zone modeling program that calculates flow, pressure, and contaminant concentrations in heating, ventilating, and air conditioning (HVAC) systems. The output of the simulation is the concentration over time of the agent at the air vents.

Air is supplied to the arena by an HVAC system that includes eight separate air handling systems (AHSs) in four mechanical rooms located at each corner of the arena, with two AHS systems in each. Each AHS system supplies ducts running in the N-S or E-W directions. The AHSs for the N-S direction have a 35,000 cubic feet per minute (cfm) capacity and supply 16 vents, while those for the E-W direction have a 20,000 cfm capacity and supply eight vents (Figure 2). Air is returned to the mechanical rooms via an air return system located at ground level in the N-S direction. Each mechanical room acts as a plenum, bringing fresh air in from the outside to mix with return air before it is recycled through the HVAC system. The system is designed to provide a mix of 20% fresh air and 80% return air to the arena.




Figure 2. Layout of Ductwork and Vents

In the scenario, it is assumed that the smallpox is released in one mechanical room in aerosol form. The aerosol consists of a one-liter suspension containing 5×10^7 cell-infecting units (CIUs) per mL, with each CIU having a mass of 5×10^{-15} g. This yields an initial concentration of 2.5×10^{-4} g/L. In the scenario, this concentration is released five minutes after the concert begins at a rate of 8.3333×10^{-7} g/s for 300 seconds. It is assumed that any filters associated with the AHSs in the mechanical room have been dismantled by the individual responsible for the release.

Table 1 provides the velocity of the smallpox agent at each vent, along with the time it takes for the agent to travel to that vent. The data presented is for the system in a steady state, and therefore the velocity at each vent is constant and the travel time can be computed based on the length of the ductwork. As the concentration of smallpox at each vent varies with time, it is not practical to include here the results for all 24 vents. However, the sum of the concentration at all vents over

Table 1. Smallpox Velocity & Travel Time at Vents

Vent #	Velocity (m/s)	Time (s)
1	7.76	0.96
2-8	3.96	1.75, 2.40, 2.93, 3.34, 3.66, 3.87, 3.98
9-16	5.08	1.75, 2.40, 2.93, 3.34, 3.66, 3.87, 3.98
17-20	3.17	1.17, 1.92, 2.41, 2.65
21-24	4.57	1.17, 1.92, 2.41, 2.65

time is provided in Figure 3. It can be observed from this figure that it takes approximately 30 minutes for all of the smallpox to be vented into the arena from the time it is released in the mechanical room. This would occur around 35 minutes after the concert begins, thereby exposing concert-goers to the full quantity of the smallpox over the duration of the concert event.

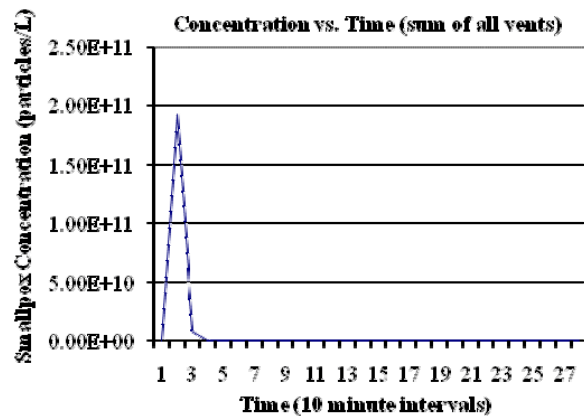


Figure 3. Concentration of Smallpox at All Vents

Airflow Computation and Agent Transport

The computation of airflow and transport of the smallpox in the arena were performed by researchers at UAB. The steps involved include the preparation of the arena geometry, discretization of the spatial domain, solution of appropriate governing equations, and post-processing of the computed results. The airflow computation requires a detailed internal geometry of the arena to form a watertight

computational domain. Because only a few 2D sketches for the arena were available, a 3D computer-aided design (CAD) model of the inside of the arena had to be reconstructed. A commercial CAD tool, Pro-Engineer (Pro/E), was used to reconstruct the design and produce a watertight geometry for mesh generation purposes. An overall view is shown in Figure 4. From this 3D CAD model, an unstructured mesh was generated by discretizing the interior of the arena. The mesh generation algorithms (Ito et al, 2006, 2007) were developed at the Enabling Technology Laboratory at UAB. The final mesh consisted of six million tetrahedral elements and 1.16 million nodes. A mesh resolution of one point every five inches was used near the seats. Two different views are given in Figure 5.




Figure 4. Reconstructed CAD Model of the Arena

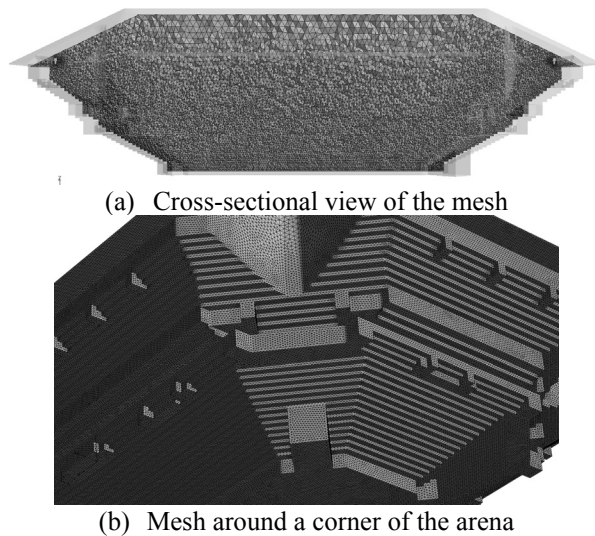


Figure 5. Mesh Used for the Air Flow and Particle Transport Inside the Arena

The speed of the airflow inside an arena is very slow and falls well within incompressible flow assumptions. The integral form of the artificial incompressibility formulation of Euler equations is taken as the governing equation for the airflow (Koomullil and

Soni, 2001). This formulation converts the elliptic nature of the governing equations to a hyperbolic form and allows the use of numerical schemes developed for hyperbolic equations to solve these equations. The spatial discretization of the governing equations uses a cell-centered, finite-volume upwind scheme (Koomullil and Soni, 1999, 2001). The convective fluxes at the cell faces are evaluated using Roe's approximate Riemann solver. Higher-order accuracy in the spatial domain is achieved using a linear reconstruction of the flow variables using the cell-averaged values.

A limiter function is also added to the linear reconstruction to avoid the creation of local extrema during the extrapolation of flow variables to the cell-faces, for higher order spatial accuracy. To handle a large number of elements in the computational domain, the flow solver is developed in a parallel framework. The decomposition of the physical domain into a set of smaller regions is achieved using METIS (Karypis and Kumar, 1995), which utilizes the graph of the grid to perform the decomposition. The Message Passing Interface (MPI) is used to pass the information across the block interfaces. A few of the computed streamlines emanating from the HVAC outlets at one corner of the arena are shown in Figure 6, showing a very complex flow pattern of air inside the arena.




Figure 6. Streamlines of Air Flow in One Corner

The smallpox is assumed to be suspended in atomized water vapor. This transport can be simulated with different levels of detail: 1) water vapor is assumed to trace the same path as the air particles; 2) water vapor is treated as a continuum and propagation of vapor is modeled using the transport equation; 3) water vapor is tracked using a Lagrangian approach; or 4) water vapor is tracked using a Lagrangian approach and a momentum source is added to the Navier-Stokes equations in a fully coupled fashion. In this effort, the third approach was used, as described below.

For the Lagrangian particle tracking, the rate of change of position and velocity of a particle can be written as

$$\frac{d\vec{X}_p}{dt} = \vec{V}_p \quad \frac{d\vec{V}_p}{dt} = \frac{\vec{V}_c - \vec{V}_p}{t_d} + \vec{g}$$

where t_d is the particle dynamic relaxation time and is defined as $\frac{4\rho_p d_p}{3C_d \rho_c |\vec{V}_c - \vec{V}_p|}$, $C_d = \frac{24}{\text{Re}} [1 + 0.15 \text{Re}^{0.687}]$,

\vec{X}_p , \vec{V}_p , ρ_p , and d_p are the position, velocity, density, and diameter of water vapor particles, respectively, \vec{V}_c and ρ_c are the velocity and density of air, respectively, \vec{g} is the acceleration due to gravity, and t is the time. These equations can be simultaneously numerically integrated to track the water particles. Because the simulations have been carried out using unstructured meshes, sophisticated search algorithms are required to locate the cell in which the particle lies as it leaves a cell. A search algorithm based on the cell-face normal is used for this.

As the simulation progresses, the number of particles accumulated on each boundary surface is counted. From this information, the amount of smallpox on each boundary face can be calculated based on the particle size distribution and number of particles per packet. In post-processing, the arrangement of the seats is mapped onto the boundary surfaces, and the amount of smallpox deposited on each seat is calculated. The potency of each agent particle is estimated based on the time the particles are exposed to air (Harper, 1961). In this study, 20 active infectious units of agent particles are taken as the critical amount for infection (Alibek, 2004). A map of the infected seats for one and two hours after release is shown in Figure 7.




Figure 7. Infected Seats at Different Times after the Release of the Aerosol

Generation of Population and Initial Infections

The generation of the population of approximately one million people and the determination of which were initially infected with smallpox at the concert were performed by researchers at JHU/APL. Both U.S. Census Bureau data for the Baltimore area and demographic data obtained from the Baltimore Metropolitan Council (BMC) were used to synthesize

the population. Each Transportation Analysis Zone (TAZ) used in the Traffic Flow simulation (described later) was considered a “Community” in the language of Longini et al, 2007. The number of households in each TAZ was taken directly from the BMC data. For the purposes of this initial prototype, “Institutions” in each TAZ (e.g., college dormitories, retirement communities) were not considered.

Individual household sizes, from one to seven, were determined based on average household sizes that were computed from the BMC data, by applying a combination of a truncated Poisson distribution and average 2005-07 Census Bureau household size distribution for the Baltimore region. This synthesized household size distribution resulted in a total population within 0.6% of the actual total population in the BMC data.

Each person was assigned an age group based on U.S. Census data for the state of Maryland. Each household was assigned at least one adult, with the rest of the household size assigned randomly based on the age group distribution. Households were grouped into Social Groups with a mean of four households and a standard deviation of one household, to approximate the (fixed) Social Group size of four used in Longini et al, 2007. Social Groups were grouped into Neighborhoods of a size as close to 500 persons as was possible, based on the number of persons in each TAZ.

The execution of the agent transport simulation described in the previous section resulted in a total of 2809 seats in the arena receiving a cumulative dosage above the infection threshold. For the purposes of the initial execution of the prototype simulation, individuals of age five and above were selected randomly from the synthetic population as having occupied those seats. (Future excursions will consider potential correlations of seat price to median TAZ income, and the likelihood of persons from the same household occupying adjacent seats.) Based on the ages of the randomly selected individuals in the infected seat locations, 10% of those born before 1971 were assumed to have residual immunity, and thus were not considered to be initially infected.

FASTER-THAN-REAL-TIME SIMULATION FEDERATION COMPONENTS AND RESULTS

Population Disease Spread and Decision-Making

The spread of the disease through the population and the population’s reaction to communications were

modeled by researchers at JHU/APL. At the core of disease spread modeling are the values used to calculate the exposure and infection of individuals. To do this requires, at every given location at every time step: the number of infected individuals and their current state of disease, and the number of susceptible individuals and their level of susceptibility (i.e., based on their age, whether they have been previously vaccinated, and whether they are taking preventive measures). Because the disease spread calculation requires the state of the entire population, the decision was made to have the disease spread simulation handle the population decisions and actions as well.

The resulting simulation is called Sparrows, which is an agent-based simulation that uses the Terracotta clustering technology to handle the large amount of calculations needed in modeling an entire city (Terracotta web site). Sparrows was built using the Master-Worker design pattern, which involves a master thread filling a work queue with tasks and worker threads who monitor the queue, removing tasks and executing the work when they become available. This design allows the number of worker threads to scale to the amount of available processors. For the executions of the federation performed in this prototype effort, Sparrows ran on a single workstation with two quad-core CPUs with eight gigabytes of RAM. As the simulation was written on Terracotta, Sparrows could be extended to run using any number of additional workstations with minimal effort. A user interface was built to allow users to examine the current status of the population and visualize the spread of the disease through the city (Figures 8(a) and 8(b)).

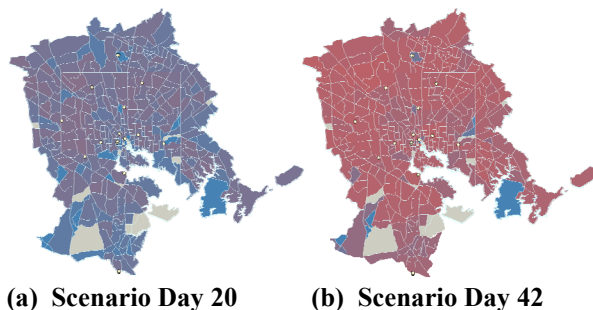


Figure 8. The TAZ regions are colored based upon the percentage of their population that is infected. The gray zones have no households. The yellow pins are the hospitals that were modeled.

The locations in Sparrows are a hierarchy of places where individuals can be exposed. At the lowest level are the Homes (mapping one-to-one to each household), Schools, Work Groups and Hospitals. Homes are combined into small groups called Social

Groups, and Work Groups are combined to form Workplaces. Social Groups, Schools, Workplaces, and Hospitals are grouped together into Neighborhoods, which are all part of Communities. In Sparrows, each Community is attached to a single TAZ.

The agents in Sparrows are the households. Operating on a diurnal cycle, each household decides what actions each member of that household will take during the next day. In a normal case, this involves each household member proceeding to his/her daily commute. Each household member has a normal commute destination: a School, a Work Group, a Hospital or the household's Home. A normal commute would then involve an individual residing at their commute destination as well as all the locations in their household's Home hierarchy (i.e., the Home, Social Group, Neighborhood and Community) with a possibility of exposure and infection at each.

These two exposure locations (commute destination and home) are mapped in Sparrows as day and night in its diurnal cycle. Individuals can infect/be infected by all other individuals at each location, with different probabilities at each (i.e., a very high probability in an elementary School versus a very low probability at the Community level). This approach was derived from the model community described in Longini et al, 2007. Individuals with the Hospital as their daily commute destination are a special case, as they represent the available staff at each hospital for the Hospital Surge simulation on any given day. As home and commute destinations are children of parent TAZs, the daily commute information between TAZs is reflected to the Traffic Flow simulation.

At each cycle, households can change the travel behavior of their members. They can change their destination to a Hospital if they feel an individual needs to seek treatment (or vaccination). They can also self-quarantine and spend both the daytime and evening in their household. Interoperating with the Hospital Surge simulation, individuals entered into Hospitals for care remain there until the Hospital Surge simulation rejects them from entering or discharges them from care.

The decision-making process for households in Sparrows is driven by the risk communications sent from the Command and Control simulation. Every day, individuals react to the most recent risk communications. Based upon the communications, their own disease state (e.g., vaccinated, infected, contagious) and that of their household (e.g., at-risk due to infected household members), individuals may

correctly self-quarantine or seek vaccination, or may exhibit counterproductive behaviors such as queuing at medical facilities as “worried well.”

Command and Control

Command and Control (C²) was modeled by researchers at JHU/APL as a constructive simulation component. The C² simulation supports analysis of strategic, tactical and operational planning and activities, including deployment of resources, maintenance of situational awareness, and public communications. DHS has released, and continues to update, a set of documents that provide guidance for emergency response agencies. The National Response Framework (NRF) (DHS 2008) establishes a single, comprehensive approach to domestic incident management to prevent, prepare for, respond to, and recover from terrorist attacks, major disasters, and other emergencies. The C² simulation follows the NRF guidelines and other federal guidance documents.

The C² simulation exchanges data with other simulations in the federation and provides initialization and shutdown services for the federation. Figure 9 shows the situational awareness dashboard maintained by the C² simulation. The simulation subscribes to data published by both Sparrows and the Hospital Surge simulation. Daily data is received for each school and each hospital. The data is processed to plot the number of new cases and fatalities, cumulative student and teacher absences from all schools, medical staff and bed shortages, and the number of people vaccinated each day.

The C² simulation publishes data to represent daily public communications. This data includes parameters to enable the Sparrows simulation to simulate the behavioral effects of the communications. The C² simulation models the intensity and frequency of risk

communications based on an internally maintained alert level. At the first indication of the crisis (day 1 in Figure 9), an order to vaccinate medical staff is published. The alert level is raised based on rules that can include factors for the rate of new cases, fatalities, school absences and vaccination facility capacity utilization. As the alert level rises, the C² simulation can issue an order to close schools (day 20 in Figure 9). The scenario used for Figure 9 raises the alert level when the five-day moving average of new cases increases more than 20%. Each time the level is raised, the simulated effectiveness of the risk communications is increased. The messages include instructions for the infected, at-risk, normal and recovered members of the population. The preferred behavior is for the infected to be isolated, the at-risk to seek vaccination, and the normal and recovered to continue with their daily routine of attending school or going to work.

The results for the scenario plotted in Figure 9 indicate that by using the five-day moving average of new daily cases, the alert level is raised quickly in anticipation of an increase in severity of the crisis – it is raised just as the number of cases begins its most rapid increase.

Traffic Flow

The modeling of traffic flow was performed by researchers at FAU using the meso/micro-simulation software Aimsun NG. The ability to simulate real-time, dynamic traffic assignment scenarios is an effective tool in locating and eliminating conflict points (areas with high density and high delay-time). The ability to visualize the network and the loading on the infrastructure provides significant knowledge for planning purposes. In this case, it assisted in the selection of routes to emergency response center locations. This was made possible by implementing a hybrid simulation technique, which incorporates two levels of simulation (meso and micro) that are meshed

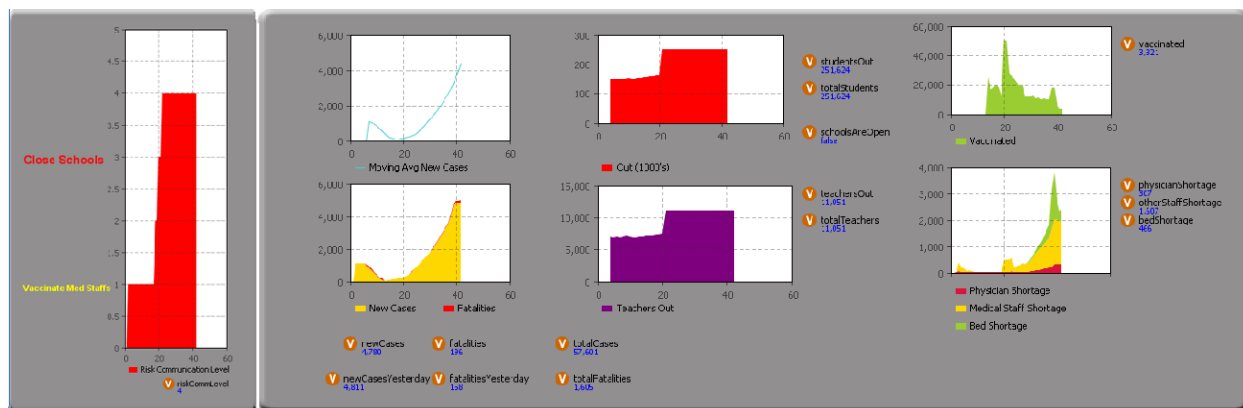


Figure 9. C² Simulation Dashboard

for enhanced analysis. The hybrid allows for simulation of small networks or areas (micro) while simulating the surrounding area (through meso). With the hybrid simulation, one can study the effects of signal controls, network origin-destination (O/D) travel routes/times and time-varying flows. Since both levels of simulation are using dynamic traffic assignment (DTA), faster-than-real-time system degradation can occur as the network demands are changing. The basis of this model involves travelers moving from one point to another through an O/D matrix.

To test the generated demand, a user equilibrium model was used, which suggests that travelers will minimize their travel time based on reviewing shortest travel times through different links in the system. Routing and re-routing throughout the network is performed with this strategy. The model uses per-link cost-based functions to determine the paths or links that will be taken based on the level of traffic demand. Therefore, DTA enables the description of traffic flow pattern evolution throughout the network. This simulation process is expressed in these steps:

1. Calculate initial shortest routes for each O/D pair, Home-Work and Home-School and Hospital trips.
2. Simulate for peak-hour trip distribution, assigning to the available routes the fraction of the trips between each O/D pair for that time interval according to the model. Obtain new average link travel times as a result of the simulation.
3. Recalculate shortest routes, taking into account the current average link travel times.
4. Relay information from step 3 back to the simulation for dynamic re-routing.

These steps were performed for 43 days of peak-hour simulation – nearly 8,000 minutes of run-time. The normal daily peak-hour trip distributions for day 1 were obtained from the Baltimore Metropolitan Council. After the disease outbreak is diagnosed, the trip levels change daily for work, school and hospital demand based on what policies the C² simulation has implemented relating to work/school closings, self-quarantine directives, and vaccination sites.

Results from the daily simulations performed as part of the initial prototype development indicated that the flow and delay times throughout the network did not increase with hospital demand. The delay time for the micro networks, on average, was nearly two times that for the meso level. Figure 10 depicts the typical flow in micro analysis to and around the hospitals after the second week. Figure 11 depicts the flow through the meso level. It differs from the micro level in that the vehicles enter the system very quickly and find their

DTA route to the micro level without significant delay. Once the vehicles enter the micro level, the delays are more pronounced due to increased inputs (signal control, car following algorithms, etc.). Overall, the traffic simulation results reinforce the importance that C² can have on limiting delays and back-ups.




Figure 10: Flow vs. Time, Micro Level




Figure 11: Flow vs. Time, Meso Level

Hospital Surge

The hospital surge simulation (HSS) dynamically predicts the hospital capacity required based on the number of casualties, individuals requiring vaccination, and worried well. In this context, hospital capacity includes beds in each of three types of units: the emergency department (ED), the intensive care unit (ICU), and the floor, as well as two types of personnel: physicians and other medical personnel. This version of the model does not include other expendable hospital supplies or non-medical personnel. The HSS models all 16 hospitals in the selected area that have licensed acute care beds necessary to treat smallpox. This includes 15 civilian hospitals and the Baltimore Veterans Administration (VA) Medical Center. Each hospital is initialized with its number of beds per unit (MHCC, 2008; Griffin, 2009), which does not change throughout the scenario.

Rather than develop an entirely new hospital model, we reviewed more than a dozen extant models. There were two limiting factors: no single model adequately modeled both treatment and vaccination as our scenario required, and most hospital surge capacity

models are focused on strategic planning or preparedness (predicting the resources required over the entire course of an incident based on a static set of inputs) and on the “pre-incident” end of the preparedness continuum. The implicit assumption is that medical and public health responses based on situational awareness do not impact the evolution of the incident, i.e., the extant models don’t support tactical planning and preparedness.

Our solutions to these problems are:

1. We selected two separate extant models, AHRQ and MaxiVac, and reverse-engineered their mathematical models into a single HSS simulation federate. The results of this reverse engineering are hospital personnel-to-patient ratios for each of the hospital units for admitted individuals (derived from AHRQ), and limits on the number of patients that can be seen based on the clinic size and available personnel (based on MaxiVac).
2. The HSS calculates the resources needed on a daily basis, allowing dynamic allocation and reallocation of medical personnel. Unlike extant models, it models the possibility that not all individuals seeking treatment will receive it in a timely manner due to hospital overload. It was for this reason that we had to reverse-engineer the mathematical models of AHRQ and MaxiVac.

Figure 12 shows the flow of individuals through the hospital. New individuals are presented to the HSS at the beginning of each diurnal cycle. The HSS models both vaccination and treatment, first passing all individuals presented to it through a triage/vaccination process. Staff requiring vaccination are “frontloaded,” i.e., they go to the front of the line because they can’t work if they’re infected, potentially impacting the hospital’s ability to treat other patients. Those individuals sent to the hospital for vaccination are vaccinated if possible based on the constraints described below. Infected individuals are admitted, if hospital capacity exists for them, and moved through the various hospital units as their disease progresses.

Individuals can be rejected in this model, unlike other models. This can happen at two points:

- Individuals can be rejected “at the hospital door,” i.e., they never get to triage/vaccination because there are too many individuals in line.
- Individuals can be identified as infected and needing treatment, but there are simply no available beds in the hospital, so they can’t be admitted. This is one area where more data is required to determine how such situations would be handled in an actual bioterrorism crisis.

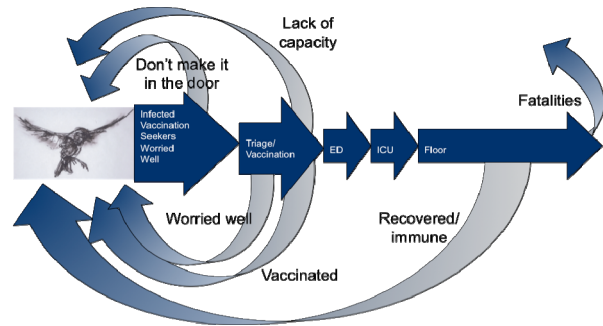


Figure 12. Patient Flow through the Hospital

Note that both categories of rejected individuals are highly likely to attempt to return to the hospital on subsequent days seeking treatment.

This model differs from other hospital surge models because it models the hospital dynamically, calculating resource requirements on a daily basis rather than projecting the number of casualties for the entire event and assuming that sufficient resources will be available. Staff members are assigned on a daily basis, allowing the modeling of an influx of staff from other sources as well as the impact of staff absences due to infection in their own household. Staff assignment is preferentially given to the triage/vaccination function because it is the bottleneck to hospital admission.

At the end of each diurnal cycle, the HSS reports hospital status to the C² simulation to support decisions on the need for more vaccination and treatment clinics and personnel. The HSS also returns recovered (immune) individuals to the population and removes fatalities, recovering renewable hospital resources such as beds and staff. Finally, it reports those who are rejected, either because they are worried well (not infected), or because the hospital has insufficient resources to treat them. This last capability is unique among hospital surge models as explained below.

Figure 13 shows the total number of patients in all the hospitals across the duration of the simulation federation prototype execution. The total number of ED beds is 650. There are 631 ICU beds and 3803 floor beds, for a total of 5084 beds available. Despite the fact that infected individuals are turned away, the total capacity of all the hospitals is never achieved. One reason for this is that the assumption that all patients must first be admitted to the ED creates a bottleneck. There are more than five times as many floor beds as there are ED beds, but patients must first be admitted to an ED bed. This suggests that investigations into altered standards of care should look at admitting to beds other than just in ED.




Figure 13. Total Hospital Occupancy

Secondly, we might still expect to see capacity achieved at some point in the scenario. Individual hospitals do achieve (or nearly achieve) capacity on different days, but the spread of the disease in different parts of the city cause this to happen at individual hospitals on different days, preventing the system of hospitals to achieve capacity on any single day.

With this framework in place, we can easily add modeling of other policies/processes, e.g., reverse triage, altered standards of care, different assignments of staff and patient transfers between hospital units.

SUMMARY AND POTENTIAL APPLICATIONS

This paper has provided an overview of a Bioterrorism Crisis Management simulation prototype, and the results of its initial execution that was performed in May 2009. As the simulation was being designed, it became evident that there were many decisions that needed to be modeled, the protocols for which have not yet been firmly established. Similarly, the probabilities of disease transmission among various segments of the population might be significantly refined based on new research in the scientific community. For these reasons, significant effort was devoted to making such factors parameters or input data for the simulation, so that they could be easily changed.

Only a few executions of the prototype simulation have thus far been performed to demonstrate its performance. One potential application of the simulation is in mission rehearsal, to study in more detail the effects that different command and control actions, and their timing, could have on the containment of the disease outbreak. Initial simulation excursion runs indicate that the degree of compliance of the population with timely and accurate directions may have significant effects on the progress of the

disease, so means of ensuring greater compliance are deserving of further study.

REFERENCES

- Alibek, K. (2004). Smallpox: a disease and a weapon, *International Journal of Infectious Diseases* 2004, 8S2:S3—S8.
- Coolahan, J.E., Kane, M.T., Koomullil, R.P., Kaisar, E.I., & Walsh, K.K. (2008). Results of a Federated Simulation of Urban Chemical Disaster Response. *Proc., 2008 Interservice/Industry Training, Simulation & Education Conference*, Orlando, FL. Department of Homeland Security (2008). *National Response Framework*.
- Griffin, K.N. (2009). Personal communication, Program Manager, Public Relations & Community Outreach (BT/135), VA Maryland Health Care System.
- Harper, G.J. (1961). Airborne micro-organisms: survival test with four viruses, *J Hyg.1961*; 59:479-486.
- Ito, Y., Shum, P.C., Shih, A.M., Soni, B.K., and Nakahashi, K. (2006), Robust Generation of High-Quality Unstructured Meshes on Realistic Biomedical Geometry, *International Journal for Numerical Methods in Engineering*, Vol. 65, Issue 6, pp. 943-973.
- Ito, Y., Shih, A.M., Soni, B.K., and Nakahashi, K. (2007), Multiple Marching Direction Approach to Generate High Quality Hybrid Meshes, *AIAA Journal*, Vol. 45, No. 1, pp. 162-167.
- Karypis, G., and Kumar, V. (1995). *METIS: Unstructured Graph Partitioning and Sparse Matrix Reordering System, Version 2.0*. Department of Computer Science, University of Minnesota.
- Koomullil, R.P., and Soni, B.K. (1999), Flow Simulation Using Generalized Static and Dynamics Grids,” *AIAA Journal*, Vol. 37, No. 12, pp. 1551-1557.
- Koomullil, R.P., and Soni, B.K. (2001), Wind Field Simulations in Urban Area, AIAA-2001-2621, Available on CD ROM, 15th AIAA Computational Fluid Dynamics Conference, Anaheim, CA.
- Longini, I.M., Halloran, M.E., Nizam, A., Yang, Y., Xu, S., Burke, D.S., Cummings, D.A.T., and Epstein, J.M. (2007). Containing a Large Bioterrorist Smallpox Attack: a Computer Simulation Approach, *International Journal of Infectious Diseases*, 11, 98-108.
- Maryland Health Care Commission (2008). *Annual Report on Selected Maryland Acute Care Hospital Services*.
- Terracotta web site, <http://terracotta.org/>.