

BioWar: A City-Scale Multi-Agent Network Model of Weaponized Biological Attacks

CASOS Technical Report

Kathleen M. Carley, Neal Altman, Boris Kaminsky,
Démian Nave and Alex Yahja

January 2004
CMU-ISRI-04-101

Carnegie Mellon University
School of Computer Science
ISRI - Institute for Software Research International
CASOS - Center for Computational Analysis of Social and Organizational Systems

Abstract

BioWar is scalable city-wide simulation, capable of simultaneously simulating the impact of background diseases, natural outbreaks and bioterrorism attacks on the population's behavior within a city. The multi-agent simulator includes social and institutional networks, weather and climate conditions, and the physical, economical, technological, communication, health, and governmental infrastructures which modulate disease outbreaks and individual behavior. Individual behaviors include health seeking, entertainment and work/school behavior. A wide variety of reports are generated based on user needs including absenteeism patterns, pharmaceutical purchases, doctor's office insurance claims reports, and hospital/emergency room reports. Sub-reports are available for specific sentinel groups including military personnel, first responders and health workers. Reports matching real world data streams and reports can be created for analyst or public health personnel including appropriate delays in generating said reports. This paper provides an overview of BioWar's current capabilities and information on the algorithms and data used to drive the simulation as of the Challenge 5 (C5) version.

This work was supported in part by Contract 290-00-0009 for Scalable BioSurveillance Technology from the Defense Advanced Research Projects Agency (DARPA) Bio-Alirt Program for work on Scalable BioSurveillance Systems, National Science Foundation (NSF) grant IGERT9972762 to the Carnegie Mellon Center for Computational Analysis of Social and Organizational Systems (CASOS), Cooperative Agreements Number U90/CCU318753-01 and UP0/CCU318753-02 from the Centers for Disease Control and Prevention (CDC), the MacArthur Foundation, the Agency for Healthcare Research and Quality and by the Carnegie Mellon Center on Computational Analysis of Social and Organizational Systems. Computations were performed on the National Science Foundation TeraScale Computing System at the Pittsburgh Supercomputer Center. Any opinions, findings, conclusions or recommendations expressed in this report are those of the authors and do not necessarily reflect the views of DARPA, the National Science Foundation, the CDC, the MacArthur Foundation, the Agency for Healthcare Research and Quality or the US Government.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JAN 2004		2. REPORT TYPE		3. DATES COVERED 00-01-2004 to 00-01-2004	
4. TITLE AND SUBTITLE BioWar: A City-Scale Multi-Agent Network Model of Weaponized Biological Attacks				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,School of Computer Science,Pittsburgh,PA,15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 38	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Keywords: bioterrorism, multi-agent modeling, social network, dynamic network analysis, syndromic surveillance, behavioral surveillance.

Table of Contents

List of Figures	ii
List of Tables	ii
1 BioWar Overview	1
2 Example of BioWar in Use	3
3 Related Work in Disease and Biological Warfare Simulation.....	3
4 BioWar Algorithms.....	6
4.1 Environment.....	6
4.1.1 Time and Ticks	7
4.1.2 Climate and Weather.....	7
4.1.3 Work and Holiday Cycle	8
4.2 Agent Population	9
4.3 Agent Activities	9
4.3.1 Daily Cycle	9
4.3.2 Agent Interaction	10
4.4 Disease	11
4.4.1 Background and Attack Diseases.....	12
4.4.2 Attacks	14
4.4.3 Transmission	14
4.4.4 Progression.....	15
4.4.5 Diagnosis and Treatment	16
5 Creating a BioWar Simulation Environment.....	17
5.1 Metropolitan Areas	17
5.2 Raw Data Sources	17
5.2.1 City Specific Data	19
5.2.2 Population Specific Data.....	19
5.2.3 Universal Data	20
5.3 Creating a BioWar Ready Input Deck	20
6 Running BioWar Simulations	20
6.1 Simulation Parameters	20
6.2 Instantiating a Simulation Population with Gensim	21
6.3 Running the BioWar Simulator	21
6.4 Computational Resources Required for BioWar	21
7 BioWar Outputs and Reports	22
8 Simulation Validation	24
8.1 Grounding Social Networks.....	25
8.2 Grounding Weather Models.....	25
8.3 Checking Simulator Outputs.....	25
8.4 Docking.....	27
8.5 Future Validation: Automation	28
Bibliography	31
Index	33

List of Figures

Figure 1: BioWar used in Detection Algorithm Development.	3
Figure 2: Validation of Simulated Average Monthly Temperatures for Norfolk, Virginia against 2001 Historical Data.	7
Figure 3: Validation of Wind Direction Frequency Distribution for SanDiego, CA against Historical Average of 1990 -1992 Data.	8
Figure 4: Process of Docking BioWar with the IPF (Incubation-Prodromal-Fulminant) Model.	28
Figure 5: Wizer Conceptual Diagram: Closed-Loop of Simulation and Inference as Experimentation.	29
Figure 6: Wizer Alert Model.	30

List of Tables

Table 1: Data Output from BioWar	2
Table 2: Additional BioWar Information Sources.....	2
Table 3: Comparison of Selected Exposure and Agent Simulators.....	5
Table 4: BioWar Location Types.....	10
Table 5: Social Network Sizes for Challenges 3 and 4.....	11
Table 6: BioWar Diseases.....	12
Table 7: Files in a BioWar Input Deck.	18
Table 8: BioWar Standard Reports.	23
Table 9: Validation Overview (by Challenge).....	24
Table 10: Validation and Tuning Methods Used (by Challenge).	24
Table 11: Input Stream Validation.....	26
Table 12: Percentage of Output Streams Empirically Validated.	27

1 BioWar Overview

BioWar is a computer simulation that combines computational models of social networks, communication media, and disease transmission with demographically resolved agent models, urban spatial models, weather models, and a diagnostic error model to produce a single integrated model of the impact of a bioterrorist attack on an urban area. BioWar is configured to represent real American cities by using census data, school district boundaries, and other publicly available information. Moreover, rather than just providing information on the number of infections, BioWar models the population of individual agents as they go about their lives – both the healthy and the infected. This allows the analyst to observe the repercussions of various attacks and containment policies as revealed through indicators such as absenteeism, medical web hits, medical phone calls, insurance claims, death rates, over-the-counter pharmacy purchases, and hospital visit rates, among others. Historically, BioWar has been used to test and improve detection algorithms for biological attacks using such indicators. Currently BioWar has been used to generate data to test detection routines from five different companies.

In addition, analysts can use BioWar to ask and answer “what if” questions of the form “what would happen to this city if three people returned from vacation with SARS?”, “what would be the first detectable indication of an aerial anthrax attack on an outdoor stadium during a game?”, or “given a ring-vaccination strategy for smallpox, what is the benefit of pre-vaccinating 10% of the health care workers”. BioWar is thus useful for preparedness training, intelligence planning, response analysis, detection algorithm evaluation, stakeholder communication, and public policy analysis.

Currently the system has been used to model five US metropolitan areas including Washington DC, Norfolk Virginia, Pittsburgh, San Diego and San Francisco. Each city is modeled using actual census, geographic, weather, school district, and business/entertainment location data. BioWar includes a symptom based disease model in which the symptoms displayed by an agent depend on their socio-demographic background and the progression of the disease. To date, 61 diseases have been modeled including smallpox and anthrax. BioWar also includes an agent self diagnosis and a physician diagnostic model. Agents can self diagnose on the basis of visible symptoms and so decide whether to stay home, purchase over-the counter drugs or go the doctor’s office or a hospital emergency room. Physicians diagnose on the basis of those symptoms and can run diagnostic laboratory tests. Note that diagnoses can be wrong. Biological attack models include aerosolized attacks and people-as-disease-carriers. Finally, there are a few preventive and response features that can be turned on or off depending on the analysts need – including vaccination, alert of medical personnel, general alert, and alert of agents who were known to be at the site of the known attack.

BioWar is currently implemented as a batch oriented, computationally intense set of programs. Simulated population sizes can range from a few thousand agents to several million, allowing small simulation runs on relatively modest systems. Simulations are repeatable and individual simulation parameters can be altered separately while using the same simulation base to observe the effect of variations in single parameters.

Versions of BioWar are developed to meet specific user needs and, for historical reasons, are called “challenges”. Challenges are numbered, starting with the Challenge 1 or C1 version of

BioWar. Each challenge version of BioWar represents an incremental increase in capability, including the generation of new reports as needed. The inputs and outputs vary between challenges so comparisons between BioWar versions must be made with appropriate care. Unless otherwise specified, this report describes the Challenge 5 or C5 version of BioWar.

Validation has been done with respect to weather and climate, social network, city layout, physician office and hospital visits, and the purchase of six broad categories of over the counter (OTC) drugs. Work is ongoing to create an automated validation and tuning tool and to increase level and type of validation.

Planned extensions to BioWar include increased fidelity of the disease model (e.g., increasing the number diseases to over 500) and communication model (modeling mass media and web-based information sources), first order models of additional attack detection methods such as tiger chips, water and air sensors, potential response models (such as quarantine, rapid drug disbursement (as with Cipro in the case of anthrax), and altered public information streams), and additional attack models to include water and food-borne attacks. We expect to continue to do optimization and validation as new features are added and new real data becomes available to us. Possible extensions include linking to various GIS systems, infrastructure models, modules for military bases over seas, and to various real time data feeds. In addition, additional extensions as needed for the Department of Homeland Security (DHS) will be added.

Table 1: Data Output from BioWar

BioWar Version	URL for Data
Challenge 1	http://legba.casos.ri.cmu.edu/biowar/c1
Challenge 2	http://legba.casos.ri.cmu.edu/biowar/c2
Challenge 3	http://legba.casos.ri.cmu.edu/biowar/c3
Challenge 4	http://legba.casos.ri.cmu.edu/biowar/c4
Challenge 5	http://legba.casos.ri.cmu.edu/biowar/c5
Note that each set of data represents a significant improvement in BioWar's functionality and level of validation.	

For additional and updated information see:

Table 2: Additional BioWar Information Sources

For BioWar general information: http://www.casos.cs.cmu.edu/projects/biowar/index.html
For current status and planned changes: http://www.psc.edu/~biowar/biowar_www/
For information on all Center for Computational Analysis of Social and Organizational Systems (CASOS) projects, including BioWar, use the CASOS general web site: http://www.casos.cs.cmu.edu
Additional model details, data sources, and slides for BioWar are available on the BioAlert web site – (note: this is a password protected site): http://www.casos.cs.cmu.edu/projects/biowar/bioAlert.html userID: guest Password: BioWar03


For BioWar papers:

- See citations [1]-[6] in the bibliography (page 31, below)
- See the CASOS general web site for copies of published papers.
- See the BioAlert web site for working papers.

2 Example of BioWar in Use

BioWar has been used to assist in the development of syndromic disease detection algorithms, for disease model verification and currently is being adapted for evaluating response strategies to biological attacks.

Figure 1: BioWar used in Detection Algorithm Development.



When used for detection algorithm development, BioWar formed a part of a repetitive development cycle:

- Detection algorithm developers specified types of needed surveillance data and the number and size of metropolitan areas for simulation.
- BioWar was enhanced and simulation environments were prepared. A set of simulations were run independently, some with and some without biological attacks, generating surveillance data, in the form of reports.
- Surveillance reports were distributed to algorithm developers (without information on when or if attacks had occurred).
- New requirements for surveillance data and target simulation areas were generated and the cycle repeats.

3 Related Work in Disease and Biological Warfare Simulation

A number of approaches have been used to study the possible effects of biological attacks. A constant in the research is the difficulty of obtaining real world data. The number of recent releases of biological agents has been few and limited in scope and past incidents of use in warfare are not particularly informative for current simulation needs.

One thread of development has been to create predictive models for the spread of biological material in an attack and the consequent risk of exposure to populations. Models developed for predicting the spread of radioactive materials and chemical hazards in accident situations provide a basis for this approach. The Consequences Assessment Tool Set - Joint Assessment of Catastrophic Events (CATS-JACE) [7]-[9] and National Atmospheric Release Advisory Center's (NARAC) NARAC Web and NARAC iClient tools [10] are examples of exposure simulations which include biological threats. Both use geographic information systems (GIS) to map the results to geographical and population features. These tools provide support for response and mitigation but are not intended to predict the progression of diseases once established in the population.

Another important thread of development has been to model the progression of disease in populations. This approach facilitates prediction and estimation of attack effects and the examination of mitigation and recovery strategies.

Models developed for the spread of infectious diseases in human populations can be harnessed for the predicting the effects of biological attack. Epidemiologists have used the SIR (Susceptible-Infected-Recovered) framing for modeling the course of epidemics [11]. Such models are typically implemented assuming homogeneous population mixing, without a spatial dimension, social (and network) dimension, or symptom-based behavior.

Cellular-automata models of artificial life for disease spread, such as the Brookings' individual-based computational model of smallpox epidemics [12], improve upon the differential model of SIR, allowing spatial operation and discontinuities. The geometry of cellular automata, however, does not match the spatial reality of the real world. Cellular automata tend to oversimplify disease propagation processes and are not amenable to calibration with empirical data.

System dynamics models such as Epi-Engine of CiMeRC (National Bioterrorism Civilian Medical Response Center) represent underlying social interactions with a system of mathematical equations [13]. System dynamics models capture the general trend of epidemics and feedback loops, but are less able to model the subtleties of micro- and meso-behaviors, and largely ignore the symbolic aspects of a population such as knowledge about school districts, recreational preferences, traffic regulations, and so on.

A discrete event simulation model of antibiotic distribution was used to examine post-exposure prophylaxis [14]. It provided useful insights, but did not model the social interactions and physical dimensions of disease spread and response.

Agent based simulations, such as BioWar, Measured Response and Episims, model at the level of the individual (the agent). The detail level and complexity of the model is limited only by the resources available to enhance the model and available computational resources. The disadvantages of this approach lie in the relatively higher cost per simulation run and the difficulties of verifying and adjusting the model.

Purdue University's Measured Response bioterrorism simulator is an agent based model using a "genome"-based sensor-action simulation model based on their Synthetic Environment for Analysis and Simulation (SEAS) [15]. It allows simulation of multiple connected geographical areas of differing sizes by the use of multi-level abstraction, where different scales are used simultaneously in the same simulation (for instance a city may be simulated at 100%

actual scale while the state is simulated at 10% scale). The Measured Response simulator omits the complexity of social interactions and does not model disease progression and symptoms and an individual's reactions to them.

Episims, an agent based model from Los Alamos National Labs' uses a transportation networks simulation to generate contact graphs that are assumed to be social networks [16]. By using this graph-based method, Episims does not assume homogeneous mixing of the simulated population. Instead it uses normal, non-attack day transportation patterns to simulate social networks and agent behaviors. In the current implementation of Episims, there is no feedback loop from agent behavior changes after infection and displaying symptoms to contact graphs. While Episims' graphs are dynamic, they are not driven by factors causing behavior changes of agents such as homophily and expertise seeking. Episims models disease spread by viral load, a significant factor among many influencing disease transmission (e.g., the capacity of the body's immune system and the transmission medium also influence spread).

Table 3: Comparison of Selected Exposure and Agent Simulators.

	BioWar	EpiSims	Measured Response	CATS	NARAC
Simulation Size	City	City	Multi City	Area	Area
Geography	US real	US real	Stylized	World	World
Population	US Census	Stylized	Stylized	US Census	?
Simulation Type	Agent	Adaptive Agent	Agent	Exposure	Exposure
GUI	No	Yes	Yes	Yes	Yes
Features					
Scalable	0-100%	0-100%	0-100%	N/A	N/A
Climate	Yes	No	No	Yes	Yes
Transport Network	Stylized	Yes	Stylized	GIS	GIS
Location Network	Yes	Yes	No	GIS	GIS
Social Network	Yes	Yes	No	No	No
Agent Learning	Yes	No	Dormant	No	No
Attack Types					
Threat Types Simulated	Biological Chemical	Biological	Biological	Biological Chemical Nuclear	Biological Chemical Nuclear
Release	Air Ground	Air Ground Direct	Direct	Air Ground	?

	BioWar	EpiSims	Measured Response	CATS	NARAC
Disease Model					
Disease Model	Symptom	Viral Load	Individual SIR	CHAS	?
Attack Diseases	4	1	1	3+	?
Environment Diseases	58	1	None	None	?
Simultaneous Diseases	Yes	No	No	No	?
Infection Mechanism	Social Net. + Random	Viral Load	Direct	Direct	None(?)
Treatment	Yes	No	No	No	?
Attack Responses					
Response	Alert	Panic	Gov.	Advice	Advice
Outputs					
Exposure Maps	No	No	No	Yes	Yes
Medical Case Data	Yes	No	?	No	No
Insurance Claims	Yes	No	No	No	No
OTC Drug Purchases	Yes	No	No	No	No
Infection Over Time	Yes	Yes	?	No(?)	No(?)

4 BioWar Algorithms

BioWar simulates individual human activity within a framework of cultural, biological and natural environmental factors. Agents move between simulated locations according to the time of day, day of the week, holiday cycle, climatic conditions and their physical condition, interacting with other agents as they do so. In addition, the spread of disease, diagnosis, treatment and the special case of the release of biological agents are simulated. These behaviors are controlled by a series of algorithms embodied in the BioWar code.

4.1 Environment

BioWar simulates environmental forces that strongly affect human behavior, including time and climate. The environment is customized for the cultural and geographical location of the simulated cities.

4.1.1 Time and Ticks

The basic unit of time in BioWar is the “tick” of four hours in duration. During a tick, weather and climate are set for the duration of the tick, the agent location and state is computed, reports generated and attacks resolved. Ticks are mapped to calendar days and seasons in a straightforward fashion: each day consists of six ticks with the days mapped to the standard (Gregorian) calendar.

4.1.2 Climate and Weather


The weather module in BioWar includes the climate and wind models. They provide distinct climate and wind patterns for the simulated regions.

4.1.2.1 Representation of Weather

The climate model generates temperature, pressure, and precipitation data for the whole period of a simulation. Climate templates of one year in length were created for each simulated city using data published by NOAA (National Oceanic and Atmospheric Administration) [17]. BioWar uses these templates to generate climate data for simulations any length. The generated yearly distributions of climate characteristics closely match the historical data for the simulated regions (Figure 2). Climate parameters are assumed to be uniform across the simulated region.

Figure 2: Validation of Simulated Average Monthly Temperatures for Norfolk, Virginia against 2001 Historical Data.

Solid black line – observed 2001 data, dashed red line – simulated data from BioWar.



4.1.2.2 Representation of Wind

The wind model generates wind speeds and direction for the whole period of a simulation. Wind is important at and after the moment of the attack, especially when the attack occurs outdoors and the biomaterial is dispersed through wind puff movement. We use a modified Gaussian Puff model of wind dispersion. The assumptions of the model are:

- The dispersed biomaterial is chemically stable and is not deposited to the ground.

- The lateral and vertical variations of the material concentration can both be described by Gaussian distributions, which are functions of downwind distance only.

Although in the simplest Gaussian model the wind speed is assumed to be constant at any height, our wind model calculates wind speed dependent on height.


An essential function of the wind model is to assess the Pasquill atmosphere stability category for the period of the attack. In the absence of detailed meteorological data, we assign a Pasquill atmosphere stability category based on the wind speed and time of the attack but not the sky condition, which is considered to be a reasonable approximation [18].

The wind model does not address building wake and seasonal differences. The wind direction changes at most one sector (10 degrees) from one simulation “tick” to the next. Currently the wind model assumes moderate insolation and thinly overcast cloud conditions.

The generated wind speed and direction distributions closely match the empirical data for the simulated regions published by the EPA (Environmental Protection Agency) [19].

A comparison between the simulated wind direction data for San Diego and historical 1990–1992 average data is shown in Figure 3. The relative difference between simulated and average historical frequency distribution values is less than 30%. Similar fidelity was obtained for other simulated regions.

Figure 3: Validation of Wind Direction Frequency Distribution for San Diego, CA against Historical Average of 1990 -1992 Data.
Solid black line – observed average for 1990-1992 data, dashed red line – simulated data from BioWar.



4.1.3 Work and Holiday Cycle

BioWar simulates the basic day cycle of an industrialized society: agents work and study during the day (if of appropriate age) and rest during the evening hours. Weekends are treated as rest days and a holiday calendar tracks major and minor American national holidays and the school vacation calendar. The school calendar includes the normal summer school vacation.

For simplicity, BioWar uses a single holiday calendar approximating national norms and a representative school district, rather than trying to simulate the full complexity of individual school district and regional holidays.

Severe weather also can interrupt the normal daily cycle on irregular basis, with heavy snowfalls being the most frequent cause of missed work and school days.

4.2 Agent Population

Simulated agents consist of a data structure and a set of algorithms to determine agent behavior. Agent characteristics such as age, sex and marital status are initialized to conform to the census demographics reported for the target metropolitan area. Agents also have a set of agent-to-agent connections (the social network) that defines strong social links between agents that is initialized based on social network research and a knowledge vector that helps define affinity between agents who come in contact.

More detailed information on agent creation is provided in Section 5 “Creating a BioWar Simulation Environment” starting on page 17.

4.3 Agent Activities

Agent activities include location based movement and interaction with other agents with corresponding possibility of infection. When agents visit locations as customers or are absent from their jobs and schools, they generate indicator data such as medical diagnoses, purchases of over-the-counter drugs, visits to medical information web sites and absentee reports as well as additional reports based on perfect knowledge (for example, the simulator knows with perfect certainty an agent’s health status, while an agent generates indicators based on perceivable symptoms in “deciding” if they should visit a pharmacy and what to buy while there).

4.3.1 Daily Cycle

BioWar advances on a tick by tick basis. Ticks are resolved separately, but the simulator takes the time of day, day of week and holiday schedule into account when determining agent activities for each tick. The basic daily cycle for agents starts at midnight with two ticks spent at home and resting, two ticks at work or school (if the agent is of the correct age) and two ticks spent at home but active. Agents may break the basic cycle by being absent from home, work or school due to their health, because they choose an alternative activity (broadly referred to as recreation) or for unspecified other reasons (a residual value based on historical absentee counts). On weekends and holidays, agents do not go to work or school. Note that this cycle is currently applied to all agents, although in reality some individuals of working age do not work or have unusual work schedules.

Agents are always placed in a geographical location appropriate for their selected activity (see Table 4, below). Locations originally supported only one type of agent activity (schools, for instance, supported students), but with the C5 version, BioWar supports both workers and “customers” at all locations (customers are consumers of the location’s service—students are a school’s customer in this sense). BioWar creates locations for the simulation based on actual economic census data as to type and number and distributes them geographically within the metropolitan area using location database information where available and randomly if not. Locations are nodes of agent activity, typically structures (such as schools, businesses or homes)

or places of public gathering (such as parks). Movement between locations is highly abstract; agents do not spend time in transit but are placed at the appropriate location at the start of each tick.

Table 4: BioWar Location Types.

Location	Definition
Home	Residences
Work	Work locations not assigned to any other category
School	Primary and secondary schools
Pharmacy	Pharmacies
Doctor	Doctor's offices
ER	Emergency rooms and hospitals
Stadium	Open air events
Theater	Indoor events
Store	Shopping location (excludes pharmacies)
Restaurant	Eating locations
University	Post secondary education institutions
Military	Military bases

Individual agent status is updated for active (living) agents each tick. Based on contacts with infectious agents or exposure to attack pathogens, infection may occur. An agent's internal health state is then advanced and activities appropriate to the agent's current location are resolved. This may include the generation of special indicator data, such as drug purchases or medical contacts. Then agent's target location for the next tick is calculated based on the agent's age, health, and the time of day.

4.3.2 Agent Interaction

While an individual agent's actions are largely determined independently of the other simulation agents, agents potentially interact with each other on every tick. BioWar uses two methods to select candidate agents for interaction: social network based and randomly. Once an agent is added to the interaction list, the interaction is resolved in a uniform way.

The social network represents strong ties between individuals, including family, friends, coworkers and classmates using the University of Chicago General Social Services (GSS) survey data with the addition of "schoolmate" for younger agents, a population not covered by the GSS [20]. Because the research data on social networks emphasizes relatively strong ties, the BioWar social network size (Table 5, below) is fixed at a relatively small size in relation to the total number of agents in the simulation. BioWar simulates a single metropolitan area at a time, so an agent's social network partners are artificially constrained to the agent population in the simulation.

Table 5: Social Network Sizes for Challenges 3 and 4.

All values are agent counts. Challenge 3 did not include Hampton City, San Francisco or Washington.

Social Network Size							
		Hampton City	Norfolk	Pittsburgh	San Diego	San Francisco	Washington DC
Expected value [21]	Range	6 - 97	6 - 97	6 - 97	6 - 97	6 - 97	6 - 97
	Mean	33	33	33	33	33	33
Challenge 3	Range		8 - 67	7 - 79	6 - 68		
	Mean		28	28	28		
Challenge 4	Range	3 - 108	2 - 101	3 - 108	3 - 110	4 - 110	3 - 111
	Mean	32	32	31	32	32	32

In addition to the social network, random interactions are used to simulate casual or chance contacts (for example, a fellow bus passenger). During each tick, BioWar adds agents from the full agent pool to the interaction list. BioWar can randomly select agents for interaction or bias the selection towards agents who are physically close.

The combined list of candidate interaction partners is then resolved. The probability of actual interaction is adjusted by the degree of similarity between agents, as represented by their knowledge vector. If the interaction occurs, agents can exchange knowledge and infectious diseases. After infection, the disease is resolved in the same way as diseases caused by exposure to biological attacks.

4.4 Disease

The current version of BioWar simulates 61 diseases -- 4 weaponized diseases and 57 naturally-occurring diseases -- simultaneously in a population. We use a symptom-based general disease model. Each disease has its own set of symptoms, timing of disease phases, variability in presentation based on age, gender, and race, and contagiousness. Each symptom has its own severity and progression timing. Furthermore, symptoms are assigned an “evoking strength” so that diagnoses based on symptoms will not only reflect accepted medical protocols but will also mimic the errors inherent in these protocols.

Each instance of a disease infecting an agent is individually represented and progressed through time as the agent goes about his or her daily business. Diseases can propagate through a population, the process of which is probabilistically determined by agent risk factors, the transmissibility of the disease, and the spatial and temporal proximity of uninfected agents to infected agents. Our disease model generates epidemic (or EPI) curves for both medically observed and total cases as output.

Certain demographic groups are more likely to be susceptible to particular diseases than others. These risk factors increase a person’s susceptibility to diseases through either host factors or environmental factors to which that person is exposed. For example, individuals who have contact with animals (sheep shearers, for example) are more likely to contract cutaneous anthrax than other occupations. In BioWar, risk factors are distributed a priori to individuals in the population according to demographic characteristics based on age, sex, race, and disease prevalence.

In constructing our disease model, we used historical accounts of known anthrax releases [22], documents from the October, 2001 bioterrorism attack [23], and disease knowledge bases

[24]-[26]. We have also drawn on the experience of other medical expert systems developed to assist in diagnosis to ground our disease model in well-founded medical knowledge representations [27].

4.4.1 Background and Attack Diseases

The current disease model supports three different disease groupings: attack diseases, outbreak diseases, and background diseases. These groupings are mutually exclusive to simplify disease modeling, although they need not be so in principle. Attack diseases are considered as a set of pathogens which might be released as part of a bioterrorism event. The introduction of attack diseases into a population is under the user's control---severity, disease type, and attack locations can be controlled separately.

Outbreak diseases are instantiated in a simulated population in a predetermined pattern, much like what might be expected over the normal course of a year. Although outbreak diseases can be controlled by the user, the default disease pattern we have designed would normally be acceptable except when special circumstances need to be considered. Unlike the first two groups, instantiation of diseases within the third group, background diseases, is controlled at simulation time by prevalence statistics gathered from California Department of Health data repositories. Background diseases are considered to be chronic diseases, so agents are selected to have background diseases at the start of the simulation based upon these statistics. Furthermore, background disease cases normally persist for the duration of the simulation.

Table 6: BioWar Diseases.

<i>Attack Diseases</i>
Bubonic Plague Cutaneous Anthrax Anthrax Inhalational Smallpox
<i>Outbreak Diseases</i>
Influenza Influenza Pneumonia Staphylococcal Gastroenteritis Food Poisoning
<i>Background Diseases</i>
Angina Pectoris Anxiety Neurosis Arteriolar Nephrosclerosis Arteriosclerotic Heart Disease Bacterial Pharyngitis Botulism Bronchial Asthma Bronchitis Chronic Simple Brucellosis Campylobacter Enteritis Cardiogenic Shock Acute

Chronic Fatigue Syndrome
Cutaneous Atypical Mycobacterial Infection
Depression
Diabetes Mellitus
Disseminated Intravascular Coagulation
Encephalitis Acute Viral
Fibromyalgia Syndrome
Giardiasis Intestinal
Gram Negative Pneumonia
Heat Exhaustion
Hepatitis A Acute
Herpes Simplex Encephalitis
Hypertensive Heart Disease
Hypovolemic Shock
Immune Deficiency Syndrome Acquired Aids
Infectious Mononucleosis
Malaria
Meningococcal Meningitis
Mycoplasma Pneumonia
Myocardial Infarction Acute
Obsessive Compulsive Neurosis
Plague Meningitis
Plague Pneumonia
Pneumococcal Pneumonia
Pulmonary Emphysema
Pulmonary Legionellosis
Salmonella Enterocolitis Non Typhi
Schistosomiasis Systemic
Shigellosis
Somatization Disorder Hysteria
Staphylococcal Pneumonia
Staphylococcal Scarlet Fever Toxic Shock Syndrome
Streptococcal Pharyngitis Acute
Streptococcus Pyogenes Pneumonia
Syphilis Primary
Tension Headache
Tuberculosis Chronic Pulmonary
Tuberculosis Disseminated
Tularemia
Tularemia Meningitis
Varicella Pneumonia
Viral Gastroenteritis
Viral Pharyngitis Acute Non Herpetic

4.4.2 Attacks

BioWar has a flexible attack model for both contagious and non-contagious pathogen release. The model lets attacks be varied by location, date, time of day, carrier agent (airborne, food borne, waterborne, among others), situation (inside or outside of a building), means of attack (land or airborne, spray or explosive delivery), pathogen, biomaterial mass, release height, efficiency, and number of attack locations (single point or multiple point). An example attack specification in BioWar is:

out large anthrax_inhalational 2002/7/4 22:00 25kg .1 300m 1.5km 7;

which generates a large, multi-point airborne attack at 22:00 on July 4, 2002 at an altitude of 300m, using 25kg of material for an attack at 10% efficiency and distributing 7 bombs along an attack line of 1.5 km in length. Note that the specification here concerns only with attacks using weaponized diseases. The BioWar disease model has other specifications for the naturally-occurring “background” or “outbreak” diseases.

Atmospheric dispersion modeling is usually performed in the local coordinate system with the origin of the system at the ground level at the point of emission (for the ground releases) or directly beneath the point of emission (for elevated releases). BioWar includes methods for translation between geographical, UTM, and local coordinate systems. When multiple ground or airborne releases are simulated, the origin of the local coordinate system is assigned to the location of the agent, and the total effect on the agent (the summary dosage) is calculated as a sum of dosages from individual releases.

The dosage inhaled by an agent is calculated using the following equation [28]-[29]:

$$1] \text{ Dose} = [QB][\pi u \sigma_y \sigma_z] \cdot \exp[-(1/2)(y/\sigma_y)^2] \exp[-(1/2)(H/\sigma_z)^2]$$

where Q is the source strength (e.g., number of anthrax spores); B is breathing rate (usually for light work $B = 5 \cdot 10^{-4} \text{ m}^3/\text{sec}$); u is wind speed in m/sec ; σ_y and σ_z are dispersion parameters that are functions of downwind distance x ; and H is height of the release in meters.

The attack model also includes methods for determining whether the agent is located in the downwind zone and how far the agent is from the point of the release. The meteorological conditions are assumed to persist unchanged over the wind puff travel time from source to receptor.

Solar ultraviolet rays may deactivate some pathogens, like anthrax. Spores released during daylight are assumed to be active for about 4 hours while an early evening release may keep the spores active for up to 12 hours. BioWar’s Attack Model takes this duration into account while estimating the effects of further travel of the released biomaterial and thus potentially infecting more agents.

4.4.3 Transmission

Disease transmission can occur via contact or air dispersal. When a person comes into contact with the transmission medium, disease transmission occurs with some specified probability. Disease transmission is governed by several probabilities: the agent’s susceptibility due to demographics, the disease’s “baserate” (ease of infection), the disease “transmissibility” (rate of transmission), and the agent’s immunization status and level, in the case of smallpox transmission.

Currently, disease transmission is mitigated only by immunization status---expressed symptoms have neither positive nor negative effect on the transmissibility of a disease. The “baserate” modifier is meant to accommodate these effects in a stochastic manner, although future BioWar versions will consider expressed symptoms (like eruptions or rashes) or the lack thereof in the disease transmission process.

4.4.4 Progression

Agents experiencing disease state transitions are modeled as nondeterministic automata. As past medical history affects these transitions, this is a non-Markovian model. At any time within the duration of a state, a medical intervention can occur and the state can be changed. The state of the disease also affects the medical intervention.

Each disease instance progresses through up to five phases:

1. Incubation: the period of time before the agent begins presenting symptoms due to a bacterial or viral infection.
2. Early symptomatic (prodromal): the period of time during which an infected agent may experience mild or non-descriptive symptoms. Many diseases omit this phase, or have no known or identifiable early symptomatic period.
3. Late symptomatic (manifestation): the period of time during which an infected agent may experience severe and/or disease-specific symptoms. In many diseases, this phase may not be distinct from the early symptomatic phase.
4. Communicable: the period of time during which an infected agent may infect other agents. This phase may overlap with the above phases. Noncontagious diseases do not have this phase.
5. Recovery/death: a period of time during which an infection resolves or causes death.

In the current version of BioWar, the length of each phase except recovery/death is generally determined uniformly randomly using a range provided by expert analysis. In the special case of weaponized inhalational anthrax, however, phase durations are determined by a lognormal distribution based upon the initial spore dosage received by an exposed agent. Recovery and death of an agent, when not affected by treatment, is determined by a Bernoulli process with p equal to the death rate of the disease among untreated victims (again, determined by expert analysis). The duration of dying and recovering is likewise stochastically determined.

Symptoms are important in BioWar on two levels. They motivate behavior and determine the initial diagnosis of agents entering the medical system. Agents with symptoms self-diagnose, stay home from work, visit their doctor or pharmacist, and change their patterns of interacting with others, depending on the severity of symptoms. This symptom-based disease model permits the representation of outliers and stochastic flux (not everyone with the same disease presents the same symptoms). The symptoms are assigned two different measures that influence which symptoms agents get and how that changes their behavior [27].

The first, frequency, is a qualitative measure of how frequently people with a particular disease will manifest a particular symptom. Frequency is denoted by a number between 1 and 5 that answers the question: “In patients with disease x , how often does one see symptom y ?” For example, patients with the diagnosis of anthrax will have a fever frequency of 5 – nearly all patients with anthrax will have fevers at some point in the course of their disease. Second, the

evoking strength is a qualitative measure of how frequently a doctor will associate a particular symptom with a particular disease.

Evoking strength is coded as a number between 0 and 5. It answers the question: “When you see symptom y, how often would a doctor think the cause is disease x?” For example, fever symptoms are not specific to any one disease – in our disease profile of anthrax, fever is given an evoking strength of 1. However, widened mediastinum is a more specific manifestation of anthrax – in patients who have a widened mediastinum, the diagnosis of anthrax should be considered thus the evoking strength for this is 5. Evoking strength is similar to specificity. Symptoms are present during both symptomatic phases with time-varying severity. Our current implementation of time-varying severity is a simple additive increase over time since a symptom was introduced.

4.4.5 Diagnosis and Treatment

As previously mentioned, we use a symptom-based differential diagnosis model to obtain information on the diseases infecting an agent who visits a medical facility. Our goal was not to build an error-free diagnosis model. Rather, we use differential diagnosis, as do medical doctors, which allow the possibility of initial misdiagnosis and the revision of diagnoses with additional information (e.g., lab results). We have based the model on the Internist1/QMR diagnosis model, but have augmented the results with probabilistic “switches” to help control aspects of the returned diagnosis (including rate of correct and incorrect diagnoses, and distribution of primary and secondary diagnoses by ICD9 code). As such, our model is not a true computational diagnostic tool, but serves to control the simulator's response to diseases in a simulated population.

Agents self-diagnose on the basis of visible or palpable symptoms. Medical personnel diagnose on the basis of visible symptoms and other information, which can include laboratory tests of varying accuracy (type 1 and 2 errors are possible) and report time. Due to the covert nature of weaponized biological attacks, doctors and ER personnel may or may not be anticipating the appearance of a particular bioagent, resulting in some degree of misdiagnosis. Moreover, doctors and ER personnel take time to file a report, delaying institutional realization of a bioattack.

Initial medical diagnosis is simulated based on the apparent symptoms and their evoking strengths. To determine which disease a person has, the groups of evoking strengths of symptoms associated with potential diseases are compared and the highest one is chosen as the diagnosed disease. In other words, the disease most strongly associated with the most severe set of symptoms is chosen. This produces a certain amount of inaccuracy, mimicking the real world. The diagnosis determines whether a person is treated properly or not and whether advanced tests are ordered. Subsequent diagnosis can update the primary diagnoses based on the appearance of new symptoms and on the results of diagnostic testing. Chief complaints are not necessarily the same as discharge diagnosis, which is consistent with observed hospital performance [30].

Diagnosis results in treatment or ordering an additional test if an agent is diagnosed at a doctor's office. If an agent reports directly to a hospital's emergency department, diagnosis results in treatment, tests, or an admission to the hospital. Treatment may not be immediately effective and symptoms vary in visibility and type of testing required for their detection. In the

current version of BioWar, treatment is modeled as a simple time-delayed probability of a success. Future versions will have more realistic treatment models.

5 Creating a BioWar Simulation Environment

The BioWar environment where agents live is based on open source data for selected metropolitan areas. This data is adapted for BioWar use and bundled into “input decks” for each simulated area or area subset. The input deck provides both a template which can be used to create a ready-to-run instance of the metropolitan area and the requisite detail data. The transformation step from input deck to instantiated city allows the creation of multiple specific instances of a metropolitan area which in turn can be used for repeated simulation runs from the same starting point.

5.1 Metropolitan Areas

BioWar simulates a scaled representation of a metropolitan area. In all simulations prepared to date, the basic unit used to create a BioWar input deck is the U. S. Office of Management and Budget Metropolitan Area (MA):

The general concept of a metropolitan or micropolitan statistical area is that of a core area containing a substantial population nucleus, together with adjacent communities having a high degree of social and economic integration with that core. Metropolitan and micropolitan statistical areas comprise one or more entire counties. [31]

Two types of metropolitan areas are used for BioWar [32]:

- Metropolitan Statistical Areas (MSA) – a single metropolitan area of fewer than 2.5 million people.
- Primary Metropolitan Statistical Area (PMSA) – a subdivision of a metropolitan area that contains more than 2.5 million inhabitants (the parent unit is termed a “Consolidated Metropolitan Statistical Area” (CMSA)).

In many simulation runs, subsets of the MSA or PMSA are used. In this case, specific counties (or county equivalents) that are the building blocks for the metropolitan area are selected from the defined metropolitan area and only data from those counties is used to build the input deck. In all cases, the simulation area consists of contiguous land areas, and any intervening water features (rivers or bays).

5.2 Raw Data Sources

BioWar uses data from many sources to build an input deck for a metropolitan area (Table 7, below). Several levels of data are used:

- Data applicable only to a specific metropolitan area, such as the outlines of census tracts, number of businesses, etc.
- Broadly applicable data that characterizes a given human population (often at a national level), such as the ages for school entry and graduation, frequency of doctor visits, etc.
- Universal data, such as disease progression, that applies to all human populations.

BioWar uses templates for specific metropolitan data and loads the detail information at run time while more universal information may be incorporated into the simulation algorithms themselves.

Table 7: Files in a BioWar Input Deck.

	<i>Input File</i>	<i>Summary</i>
Experiment Configuration	sim.cfg	Contains simulation “control knobs,” like number of agents, interaction modifiers, etc.
	biowar.ini, gensim.ini	Configurable paths for finding ‘biowar’ and ‘gensim’ input files
	behavior_Ts.txt	Specifies disease severity thresholds for selecting primary agent activities
	attacksdef.txt	Defines attack scenarios
	attacksizes.txt	Defines baserate and transmissivity for attack diseases to achieve small, medium, and large numbers of casualties
	strainsdef.csv	Defines disease outbreak patterns (locations, number of outbreaks, and severity)
Modifiers to Agent Behaviors	dv_by_age_injury.txt dv_by_gender_injury.txt dv_by_race_injury.txt ev_by_age_injury.txt ev_by_gender_injury.txt ev_by_race_injury.txt	Probability of doctor and ER visits by demographic profile due to injury causes. From CDC National Ambulatory Medical Care Survey (NAMCS).
	dv_by_age_nvisit.txt dv_by_gender_nvisit.txt dv_by_race_nvisit.txt ev_by_age.txt ev_by_gender.txt ev_by_race.txt	Probability of doctor and ER visits by demographic profile due to disease causes. From CDC National Ambulatory Medical Care Survey (NAMCS).
	nu_by_age_race.txt nu_by_gender_race.txt	Probability of web usage by demographic profile due to disease causes.
	sa_by_age.txt wa_by_age.txt	School and work absenteeism rates by age.
	diseases.csv injuries.csv	Disease and injury specifications (disease phase timings, death rate, incidence/prevalence). From medical expert and available literature.
Disease & Diagnosis Model	PMH.csv HPI.csv PE.csv SimpleTest.csv ModerateTest.csv ComplexTest.csv	QMR findings tables. From medical expert with reference to QMR installation.

	Input File	Summary
	HPI-to-OTC.csv	Maps HPI and PE findings (symptoms) to
	PE-to-OTC.csv	corresponding OTC drugs taken to relieve symptom.
	PE2ICD9.csv	Maps PE findings to ICD9 classification. From ICD9.
City & Environment Configuration	ct.txt ct_attr.txt sd.txt sd_attr.txt zcta5.txt zcta5_attr.txt	Cartographic boundaries for mapping Census data to geography. From Census 2000.
	gt_d00a.dat gt_d00.dat	Detailed geography for realistic distribution of agent and building locations. From Tiger/LINE 2002.
	ct_demo.csv	Census demographics by census tract. From Census 2000.
	sch_demo.csv	School demographics by school. From NCES CCD 1999-2000.
	emp_stats.csv	Job location counts and employment statistics. From Census County Business Patterns (CBP) 2000.
	wind_template.txt climate_template.csv	Weather generation inputs.

5.2.1 City Specific Data

At the city level, information is needed on the population demographics and distribution, the number and character of locations and the outlines of certain political and administrative units. Population demographics are drawn from US Census data on a census tract level. This data is used both to create individual agents and to locate agent homes. Other simulated locations are created by combining school location data, economic census data and landmark database data. School district shape data and school location information is used to place students in public schools where they reside (private schools are not simulated in C5 BioWar). County shape files give the outlines of the simulated area and define where locations can be placed. Climactic data is also defined at the city level.

City level data is collected, translated where necessary and stored as files as part of the input deck. The input deck is then instantiated for BioWar runs based on user configurable parameters such as population scale.

5.2.2 Population Specific Data

BioWar uses many data sources to create a high fidelity simulation environment and to realistically regulate agent behavior. As with any simulation of human society, much of agent behavior is regulated by cultural norms rather than reflecting the range of all possible human behavior. BioWar is constructed to reflect the demographic and cultural norms of contemporary American society. While it would be possible to adopt BioWar to any industrialized society, the simulator is does not attempt to make algorithmic behavior fully configurable through parameter files.

Population specific data includes:

- Social network partner types, interaction rates and network sizes.
- Family size and composition.
- Probability of absence and recreational preferences.
- Rates of preexisting medical conditions.
- Likelihood of seeking medical assistance, by type.
- Workday duration, workweek and holiday schedules.

For many categories, this data is further refined by the agent's age, sex and racial profile. These values may reside in parameter files or within the program source.

5.2.3 Universal Data

Some elements of the simulation are relatively or wholly consistent across all cultures and places. The progression of untreated disease, the dispersion of biological agents in the atmosphere and the characteristics of disease organisms are examples.

Generally the universal values are treated in BioWar in the same way as population specific data: values are held either in parameter files or are coded into the program source.

5.3 Creating a BioWar Ready Input Deck

BioWar input decks are prepared in the following sequence:

- Identification of data source(s).
- Data set download.
- Subset selection (if required, for a given metropolitan area for instance).
- Reformatting (if required) and rewrite to data files.
- Bundling to an input deck.

When complete, an input deck holds all parameters necessary to execute BioWar for the specified city or city subset.

6 Running BioWar Simulations

6.1 Simulation Parameters

BioWar contains two main components – the source code and the simulation environment.

Separating them allows for more flexibility. Adding more simulation regions as necessary, adding more diseases or changing the characteristics of the existing simulation environment do not require changing or recompiling of the source code. The simulation environment contains all the descriptions of the simulated cities including climate and wind patterns, population characteristics, disease databases, etc.

BioWar runs from the command line using batch style processing. Several Unix shell utilities help create a directory for the runs for a particular city and inside that directory scenario directories for the individual runs. The parameters for a run are defined in the text file “sim.cfg” that is located within the scenario directory. There are many parameters here that may be

modified including the length of the simulation, scaling factors for city and agents, activation of diagnosis and/or treatment of the attack diseases, response strategies and technical parameters that refer to the submodels within BioWar. Another important file within the scenario directory is “attacksdef.txt” that contains the full specification of the attack (see the example in the Attack Model section).

6.2 Instantiating a Simulation Population with Gensim

Compiling BioWar creates two executable files – “gensim” and “biowar”. Running “gensim” against the scenario directory creates the actual simulated city and provides full initialization before actual simulation starts. At this step many things happen. The random number generator is initialized, census tracts are loaded, the agent population is generated according to the defined scale and known demographics, a social network is created for each agent, school districts and school population are formed, workplaces, pharmacies, hospitals, doctor offices, entertainment venues (stadiums, theatres, stores, restaurants), military bases are created and populated, the disease database is loaded, climate and wind data are simulated for the entire run and attacks are created according to the description in the “attacksdef.txt” file.

6.3 Running the BioWar Simulator

Running the executable file “biowar” against the scenario directory performs the actual simulation. A short description of what happened at each simulation “tick” (calendar and time information, number of the outbreaks, number of infected/died people, number of transmitted diseases) is rendered on the screen. The data for all reports is generated and written to each report. After the run the value of the seed for the random number is preserved in the file “seed.csv”. Using it allows permits identical runs for the exactly same input population and facilitates testing the effects of code changes.

6.4 Computational Resources Required for BioWar

BioWar is currently designed for large-scale simulation, and thus requires substantial computational power to achieve meaningful results. Simulations of 275,000 agents (roughly 20% of the population of the currently available cities) take about 4.5 hours on a HP Alpha 1GHz 4-way SMP. For one-shot simulations, the computation time may be prohibitive, and interactive turnaround times are not possible. However, in many cases, large parametric searches may be required to simulate and validate a wide range of attack scenarios.

To perform these types of parametric searches, we use the computational resources of the Pittsburgh Supercomputing Center's TeraScale machine to execute hundreds of simulations simultaneously. The TeraScale machine is composed of 750 HP Alpha 1GHz 4-way SMP machines. Using this facility, a parametric search over several variables takes only slightly (roughly 15%) longer than a one-shot experiment. The resulting datasets can be analyzed using any spreadsheet or database software, and with the Wizer system, described below.

Future versions of BioWar will employ sophisticated statistical and simulation techniques to produce results for large cities while requiring simulation of a fraction of the total population. This will dramatically reduce experiment turnaround time to better support decision making processes.

7 BioWar Outputs and Reports

A single BioWar simulation produces a variety of reports to describe the simulator's behavior over time. The reports fall into two categories: "debug" reports and "challenge" reports. Debug reports give detailed information on the internal behavior of the simulator, while challenge reports describe aggregate population behavior and provide data streams to various event detection algorithms. Both debug and challenge reports can be targeted for sentinel populations, like health workers, for additional data streams.

The report development infrastructure is sophisticated, but flexible, to allow straightforward implementation of new reports. In many instances, an existing report can be duplicated to produce a new report with minimal development overhead. Because the reporting infrastructure is flexible and uniform, even entirely new special-purpose reports can be developed in a very short amount of time. Sentinel reports are derived directly from the whole-population reports, so targeting specific subpopulations is straightforward once the global report has been integrated into the code.

Table 8: BioWar Standard Reports.

	Report Name	Description
Debug Reports	activity	Location visit & agent activity counts (per tick)
	actual_incidence	Number of new cases of each disease (per tick)
	actual_prevalence	Total number of existing cases of each disease (per tick)
	actual_symptom_incidence	Number of agents with new symptoms of each disease (per tick)
	actual_symptom_prevalence	Number of agents with symptoms of each disease (per tick)
	anthrax_attack	Detailed statistics on anthrax attacks (per tick)
	avevisits	Number of pharmacy, doctor, & ER visits (summary)
	deaths	Number of deaths due to each disease (per tick)
	disease_info	Statistics on disease model & disease progression (per tick)
	dpurchase.txt	Number of each drug purchased (per tick)
	interact	Basic interaction statistics (per tick)
	interact_day	Detailed interaction statistics (per day)
	observed_incidence	Number of newly <i>diagnosed</i> cases of each disease (per tick)
	observed_prevalence	Total number of <i>diagnosed</i> cases of each disease (per tick)
	perdocvisits	Total number of doctor visits (summary)
	perervisits	Total number of ER visits (summary)
	perpharmvisits	Total number of pharmacy visits (summary)
	seed	Random seed for current experiment
	total_absenteeism	Work & school absenteeism rates (per tick)
Challenge Reports	edregistration	ER registration data (per tick)
	<u>edregistration_health_workers</u>	Health (ER) workers sentinel report
	<u>edregistration_military</u>	Military personnel sentinel report
	insuranceclaim	Doctor visit data (per tick)
	<u>insuranceclaim_health_workers</u>	Health (ER) workers sentinel report
	<u>insuranceclaim_military</u>	Military personnel sentinel report
	pharmacy	Sales of each drug per pharmacy (per tick)
	<u>pharmacy_health_workers</u>	Health (ER) workers sentinel report
	<u>pharmacy_military</u>	Military personnel sentinel report
	school	Number of agents at each work location (per tick)
	work	Number of agents at each school location (per tick)

8 Simulation Validation

We currently validate the simulator by docking (model-to-model behavior and output alignment and comparison), by grounding the social networks and weather models on empirical data, and by checking the simulator output against the bounds and means of empirical data.

Table 9: Validation Overview (by Challenge).

	C1	C2	C3	C4	C5
Number of types of validation	1	3	5	7	7
Number of types of input streams matched	3	6	12	14	15
Number of output streams matched	2	2	4	12	12
Level of match on output streams	pattern	pattern	bounds	C3 + mean, docking	C4 + by month and daily
Level of real data	1 city Pattern	1 city Pattern Peaks	C2 + national bounds	C3+OTC streams + SDI streams	C4 + additional SDI streams
Validation Indicator - Input [†]	0.061	0.139	0.483	0.794	0.892
Validation Indicator - Output [†]	0.077	0.231	0.378	0.636	
Overall Validation	0.069	0.185	0.431	0.715	
[†] Validation Indicator = (Sum across replications (Sum across all input and output data streams of (successful validation by Type * Number of Levels of validation done) / (MaxType relevant * Number of possible levels of validation relevant)) / number of data streams)) / Number of replications Type ▪ Generic Matching ▪ Characteristic Matching ▪ Relative Timing of Peaks ▪ Within Bounds ▪ First Moment ▪ Empirical Pattern ▪ Docking (not relevant for input) Levels of validation – illustrative (will vary by stream), if correlation is done then all 4 levels are done ▪ Yearly ▪ Seasonal ▪ Monthly ▪ Week Day					

Table 10: Validation and Tuning Methods Used (by Challenge).

Validation Type	C1	C2	C3	C4	C5
Docking: Comparison against another model.				■	■

Validation Type	C1	C2	C3	C4	C5
Generic Pattern: Showing that the pattern for each generated data stream matches observed patterns	■	■	■	■	■
Characteristic Matching: Showing for each generated output data stream that it has correct seasonal or daily pattern.		■	■	■	■
Relative Timing of Peaks: Showing time between peaks for different data streams matches observed difference.		■	■	■	■
Empirical Pattern: Showing pattern for each generated data stream matches the empirical pattern – best for input streams.			■	■	■
Within Bounds: Showing for each generated output data stream that the mean of simulated stream falls within min/max of that stream for real data.			■	■	■
First Moment: Showing for each generated output data stream that the mean is not statistically different than real data – yearly, monthly or daily.				■	■
■ Indicates validation attempted.					

8.1 Grounding Social Networks

The social network is built from agents in the simulation, thus ensuring that the demographics of the social network match the simulated population and are as accurate as the census data used to build to population. Certain connections are also constrained using census data: for instance family sizes are constrained to match fertility rates for American females by age, sex and marital status. These values were spot checked during code implementation but are not checked for each run.

The social network itself is checked for correct size using the Klov Dahl study [21] as the target value (see Table 5 on page 11). Degrees of separation have also been experimentally computed for some smaller simulation populations, but the algorithm in current use is too computationally intensive for regular checking of large simulation decks. The degrees of separation for the checked population is artificially low: from 3-4 while estimates for the general population range from 6-7. The lower than expected value is an artifact of the bounded population size (agents can only connect to other agents in the simulation, not the population of the world at large).

8.2 Grounding Weather Models

The climate and wind models are based on recorded weather for the target metropolitan area. The models generate weather for the simulation based on template values, but do not use the historical data directly. The output for both models is graphed against historical data and visually checked for variance (see Figure 2 on page 7 and Figure 3 on page 8 for example output).

8.3 Checking Simulator Outputs

We check the simulation outputs against the bounds and means of empirical data. The check routine takes as inputs the simulation outputs of:

1. Number of absences and registered students per day for each school (school.csv)
2. Number of absences and registered workers per day for each workplace (work.csv)

3. Number of visit records per day for each emergency room (edregistration.csv)
4. Number of visit records per day for each doctor office (insuranceclaim.csv), assuming that each visit produces one insurance claim.
5. Number of units of seven types of over the counter drugs purchased per day in each pharmacy (pharmacy.csv)

We then post-processes the above to give:

1. Absenteeism as a percentage of registered students per day across all schools
2. Absenteeism in percent per day across all workplaces
3. Number of visit per person per year across all emergency rooms
4. Number of visit per person per year across all doctor offices
5. Number of units of the seven types of drugs purchased in each pharmacy per day, per day-of-week, per month-of-year, and per year.

We then compare the post-processed output data with the empirical data in two ways:

1. Compare the maximum and minimum empirical bounds, giving an alert if the simulated output falls beyond the bounds
2. Compare the means using the Smith-Satterthwaite procedure which computes the degrees of freedom to compare means with unequal variances, checking whether the means are statistically significantly different or not.

Table 11: Input Stream Validation.

Input Stream Type	C1	C2	C3	C4	C5
Wind			■	■	■
Climate				■	■
Calendar		■	■	■	■
Disease models (symptoms, timing)			■	■	■
Disease prevalence			■	■	■
Diagnostic tests				■	■
Injury					■
Population characteristics	■	■	■	■	■
Social Network			■	■	■
School	■	■	■	■	■
Occupations			■	■	■
Activity locations (e.g., restaurants)			■	■	■
Zip codes		■	■	■	■
Behavioral differences	■	■	■	■	■
Drug-Symptom mapping		■	■	■	■
■ Indicates validation attempted.					


Table 12: Percentage of Output Streams Empirically Validated.

Metropolitan Area	C3	C4
San Francisco		5/12 (41.67%)
San Diego	3/4 (75.0%)	9/12 (75.00%)
Pittsburgh	2/4 (50.0%)	9/12 (75.00%)
Norfolk	2/4 (50.0%)	8/12 (66.67%)
Hampton City (Norfolk subset)		4/12 (50.00%)
Washington DC		4/12 (33.33%)

This checking process is the basis for the “alert” part of Wizer, an automated validation tool described in Section 8.5.

8.4 Docking


We have docked the BioWar model for anthrax against a revised SIR model for anthrax called the IPF (Incubation-Prodromal-Fulminant) model. This revision to the SIR model is necessary because anthrax is infectious but is not contagious. The IPF model distinguishes three stages of anthrax disease progression: incubation, prodromal, and fulminant. Figure 4, below, shows the process of docking BioWar with IPF. For further information see: Model Alignment of Anthrax Attack Simulations, Li-Chiou Chen, et.al. [2].

Figure 4: Process of Docking BioWar with the IPF (Incubation-Prodromal-Fulminant) Model.

8.5 Future Validation: Automation


In the near future we plan to validate BioWar using an automated validation tool called Wizer (**What-If Analyzer**). Wizer is a system that combines an inference engine and simulation virtual experiments to do what-if analyses and validation. It explores the space of possible parameters and models by performing empirical-data-driven knowledge-intensive search steps via an inference engine constrained by the simulation, instead of just doing statistical and mathematical calculations or rule-based inference. Wizer operates on both numerical and symbolic space. Wizer uses an approach which mimics a human scientist doing experiments through hypothesis, experiment design, experiment execution, data gathering, inference, and uncovering of causal relations. While the “alert” part of Wizer has been implemented, the inference engine part is still in development. Figure 5 (page 29) shows a conceptual diagram of Wizer, while Figure 6 (page 30) shows the dataflow for the “alert” part of Wizer.

Figure 5: Wizer Conceptual Diagram: Closed-Loop of Simulation and Inference as Experimentation.



As shown in Figure 5, Wizer Alert determines which data streams are wrong and how they are wrong. The Wizer Inference Engine takes the simulator’s causal diagram of what parameters influence which data and the empirical constraints and confidences on parameters to make a judgment on which parameters to change and how. This results in new parameters for the next simulation. This simulation in turn yields simulation outputs which are fed into Wizer Alert. This cycle repeats until sufficient validity is achieved based on user-defined criteria.

Figure 6: Wizer Alert Model.



As shown in Figure 6, Wizer Alert takes as input empirical data on school absence, workplace absence, doctor visits, emergency room visits, SDI (Surveillance Data Inc.) emergency room visits data and over-the-counter drug purchases data. It also inputs the outputs of BioWar simulator. It then calculates the statistics (if needed) for the simulated outputs and for the empirical data and compares the statistics. The Wizer Alert produces results by conducting minimum bound checking, maximum bound checking and mean comparison.

Bibliography

- [1] K. M. Carley, D. Fridsma, E. Casman, N. Altman, J. Chang, B. Kaminski, D. Nave, and A. Yahja. “BioWar: Scalable Multi-Agent Social and Epidemiological Simulation of Bioterrorism Events”. *NAACSOS conference proceedings*, Pittsburgh, PA, 2003.
- [2] L. Chen, K. M. Carley, D. Fridsma, B. Kaminsky, and A. Yahja, “Model Alignment of Anthrax Attack Simulations”, 2003, in review.
- [3] A. Yahja, and K. M. Carley, “Generating Realistic Heterogeneous Agents: Computing Confidant-based Base Interaction Probabilities”. *NAACSOS conference proceedings*, Pittsburgh, PA, 2003.
- [4] A. Yahja, and K. M. Carley, “WIZER: What-If Analyzer for Automated Social Model Space Exploration and Validation”. *NAACSOS conference proceedings*, Pittsburgh, PA, 2003.
- [5] A. Yahja. Wizer: Automated Validation of Large-Scale Multi-Agent Systems. CASOS Technical Report, Carnegie Mellon University, Pittsburgh, PA, 2003
- [6] A. Yahja, “BioWar: Simulating Disease Outbreaks using Social Networks”. *CASOS 2002 conference proceedings*, Pittsburgh, PA, 2002.
- [7] http://www.dtra.mil/td/acecenter/td_hpac.html, Retrieved: 10/2003.
- [8] <http://cats.saic.com>, Retrieved: 10/2003.
- [9] <http://www.saic.com/products/simulation/cats/cats.html>, Retrieved: 10/2003.
- [10] <http://narac.llnl.gov/index.html>, Retrieved: 10/2003.
- [11] R. M. Anderson and R. M. May, *Infectious Diseases of Humans: Dynamics and Control*, Oxford University Press, New York, 1991.
- [12] J. M. Epstein, et. al., “Toward a Containment Strategy for Smallpox Bioterror: An Individual-Based Computational Approach”, Brookings Institution, CSED Working Paper No. 31.
- [13] <http://www.cimerc.org/projects/default.htm>, Retrieved: 2003.
- [14] N. Hupert, A. I. Mushlin, and M. A. Callahan, “Modeling the Public Health Response to Bioterrorism: Using Discrete Event Simulation to Design Antibiotic Distribution Centers”, *Medical Decision Making*, 22:1-9, 2002.
- [15] A. R. Chaturvedi and S. R. Mehta, “Simulations in Economics and Management: Using the SEAS Simulation Environment”, *Communications of the ACM*, March 1999.
- [16] S. Eubank, “Scalable, Efficient Epidemiological Simulation”, SAC, 2002.
- [17] <http://www.ncdc.noaa.gov/oa/ncdc.html>, Retrieved: 2003.
- [18] R. Barrat, “Atmospheric dispersion modeling: an introduction to practical applications”, *Earthscan*, ISBN 1 85383 642, 2001.
- [19] <http://www.epa.gov>, Retrieved: 2003.
- [20] <http://www.icpsr.umich.edu:8080/GSS/homepage.htm>, Retrieved: 2003.
- [21] A. S. Klov Dahl, “Social Networks in Contemporary Societies”, Dynamic Social Network Modeling and Analysis workshop, Washington DC, November 7-9, 2002, *unpublished*.

- [22] T. V. Inglesby, et.al., “Anthrax as a Biological Weapon”, *Journal of American Medical Association*, Vol. 281, No. 18, May 1999.
- [23] http://www.cdc.gov/ncidod/EID/vol8no10/contents_v8n10.htm, Retrieved: 2003.
- [24] USAMRIID’s Medical Management of Biological Casualties Handbook, US Army Medical Research Institute for Infectious Diseases, 2001.
- [25] K. H. West, “Infectious Disease Handbook for Emergency Care Personnel”, ACGIH, 2001.
- [26] C. M. Isada, B. L. Kasten, M. P. Goldman, L. D. Gray, and J. A. Aberg, “Infectious Disease Handbook”, AACC, 2003.
- [27] R. A. Miller, H. E. Pople, and J. D. Myers, "Interist-I, An Experimental Computer-based Diagnostic Consultant for General Internal Medicine", *N Engl J. Med*, 07:468-76, 1982.
- [28] D. Bruce Turner, “Workbook of Atmospheric Dispersion Estimates: An Introduction to Dispersion Modeling, Second Edition”, Lewis Publishers, ISBN 1 56670 023 X, 1994.
- [29] M. Meselson, “Note Regarding Source Strength”, *ASA Newsletter*, article 01-6a (www.asanltr.com), Retrieved: 2003).
- [30] E. M. Begier, D. Sockwell, L. M. Branch, J. O. Davies-Cole, L. H. Jones, L. Edwards, J. A. Casani, and D. Blythe, “The National Capitol Region’s Emergency Department Syndromic Surveillance System: Do Chief Complaints and Discharge Diagnosis Yield Different Results?”, *Emerging Infectious Diseases*, 9(3):393-396, 2003.
- [31] <http://www.census.gov/population/www/estimates/metroarea.html>, Retrieved: 12/2003.
- [32] <http://www.census.gov/population/www/estimates/aboutmetro.html>, Retrieved: 12/2003.

Index

activities	9	Consolidated Metropolitan Statistical Area	17
additional information.....	2	daily cycle	9
agent		data sources.....	17
activities	9	death.....	15
concept	9	diagnosis	16
interaction	10	discrete event simulation model.....	4
atmospheric dispersion modeling	14	disease	11
attack diseases	12	attack.....	12
attacks	14	background.....	12
background diseases.....	12	communicable	15
Bernoulli process	15	death.....	15
BioWar		diagnosis	16
additional information.....	2	early symptomatic.....	15
agent.....	9	evoking strength.....	16
attacks	14	incubation.....	15
challenges.....	1	late symptomatic	15
climate.....	7	manifestation.....	15
computation resources	21	outbreak.....	12
data sources	17	prodromal	15
example of use	3	progression.....	15
executable files.....	21	recovery.....	15
extensions to.....	2	symptoms	15
input deck.....	17	transmission	14
output from.....	2	treatment	16
reports	22	docking.....	24
running	20	early symptomatic.....	15
tick.....	7	Epi-Engine	4
validation.....	24	Episims.....	5
versions	1	evoking strength.....	16
weather.....	7	exposure models.....	4
wind.....	7	Gaussian Puff	7
<i>biowar</i> executable file.....	21	General Social Services survey.....	10
California Department of Health data		<i>gensim</i> executable file.....	21
repositories.....	12	geographical locations	9
CATS-JACE	4	grounding	24
cellular-automata.....	4	GSS	10
challenge		holiday calendar	8
C1	1	ICD9 code	16
C5.....	2	iClient.....	4
Challenge 1	1	incubation.....	15
Challenge 5	2	Incubation-Prodromal-Fulminant	27
checking	24	input deck.....	17, 20
climate.....	7	interaction	10
CMSA	17	Internist1	16
communicable	15		

IPF	27	Smith-Satterthwaite procedure.....	26
late symptomatic	15	social network	10
MA	17	susceptibility	11
manifestation.....	15	Susceptible-Infected-Recovered	4
Measured Response	4	symptoms	15
Metropolitan Area	17	syndromic disease detection	3
Metropolitan Statistical Area	17	Synthetic Environment for Analysis and	
MSA	17	Simulation	4
NARAC iClient.....	4	TeraScale.....	21
NARAC Web	4	tick.....	7
outbreak diseases	12	transmission	14
output	2	treatment	16
parameters	20	ultraviolet rays	14
Pasquill atmosphere stability category.....	8	vacation calendar	8
PMSA.....	17	validation.....	24
Primary Metropolitan Statistical Area	17	checking	24
prodromal.....	15	docking.....	24, 27
progression.....	15	grounding	24, 25
progression models	4	weather	7
QMR	16	wind.....	7
recovery.....	15	Wizer	28
reports	22	Alert	29
SEAS	4	Inference Engine	29
SIR	4		