

Agent Modeling of a Sarin Attack in Manhattan

Venkatesh Mysore
New York University
715 Broadway #1012
New York, NY, USA
mysore@cs.nyu.edu

Giuseppe Narzisi
University of Catania
V.le A. Doria 6, 95125
Catania, Italy
narzisi@dmi.unict.it

Bud Mishra
New York University
715 Broadway #1002
New York, NY, USA
mishra@nyu.edu

ABSTRACT

In this paper, we describe the agent-based modeling (ABM), simulation and analysis of a potential Sarin gas attack at the Port Authority Bus Terminal in the island of Manhattan in New York city, USA. The streets and subways of Manhattan have been modeled as a non-planar graph. The people at the terminal are modeled as agents initially moving randomly, but with a resultant drift velocity towards their destinations, e.g., work places. Upon exposure and illness, they choose to head to one of the hospitals they are aware of. A simple variant of the *LRTA** algorithm for route computation is used to model a person's panic behavior. Information about hospital locations and current capacities are exchanged between adjacent persons, is broadcast by the hospital to persons within its premises, and is also accessible to persons with some form of radio or cellular communication device. The hospital treats all persons reaching its premises and employs a triage policy to determine who deserves medical attention, in a situation of over-crowding or shortage of resources. On-site treatment units are assumed to arrive at the scene shortly after the event. In addition, there are several probabilistic parameters describing personality traits, hospital behavior choices, on-site treatment provider actions and Sarin prognosis. The modeling and simulation were carried out in Java RePast 3.1. The result of the interaction of these 1000+ agents is analyzed by repeated simulation and parameter sweeps. Some preliminary analyses are reported here, and lead us to conclude that simulation-based analysis can be successfully combined with traditional table-top exercises (as war-games), and can be used to develop, test, evaluate and refine public health policies governing catastrophe preparedness and emergency response.

Categories and Subject Descriptors

I.6.5 [Simulation and Modeling]: Model Development—*Modeling methodologies*; I.6.3 [Simulation and Modeling]: Applications; J.4 [Social and Behavioral Sciences]: Sociology; J.3 [Life and Medical Sciences]: Health

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

General Terms

Experimentation, Security, Human Factors, Verification

Keywords

Terrorism, Emergency Response, RePast, LRTA*

1. INTRODUCTION

New York University's Center for Catastrophe Preparedness and Response (CCPR) was founded in the wake of the cataclysmic terrorist attacks on the World Trade Center in New York city. As part of its Large Scale Emergency Readiness (LaSER) project, mathematical models of the dynamics of urban catastrophes are being developed to improve preparedness and response capabilities. The need for emergency response planning has been reinforced by the recent string of natural calamities and controversies over the non-implementation of suggested plans (for example, see the hurricane Katrina disaster predicted and analyzed well-before the event [11]). Conventional policy planning relies largely on war-gaming, where the potential disaster scenario is enacted as a table-top exercise, a computer simulation or an actual full-scale rehearsal using actual resources and players. It has been repeatedly observed that "disaster planning is only as good as the assumptions on which it is based" [3].

Agent Based Modeling (ABM) is a novel technique for simulating and analyzing interaction-based scenarios [9], with its recent application to disaster management. The first scenario we investigated was the 1998 food poisoning of a gathering of over 8000 people at a priest's coronation in Minas Gerais, Brazil leading to 16 fatalities [7]. Multi-agent modeling was explored for this problem by allowing simplistic hospital and person agents to interact on a 2-dimensional integer grid. Counter-intuitive and unanticipated behaviors emerged in the extremely parameter sensitive system, immediately suggesting a potential use for such agent-simulation-based analysis of catastrophes. This paper provides a more thorough and practical example of how a large-scale urban catastrophe can be modeled, how real data about maps, subways and hospitals can be integrated, how person, hospital and on-site responder behavior can be modeled, and how simulations can be analyzed to yield tangible non-trivial inputs that a team of expert policy makers and responders can utilize, in conjunction with conventional approaches.

Specifically, we picked the nerve gas agent Sarin and the city of Manhattan to demonstrate our tools and techniques. Our choice was based on the literature available about a similar attack executed in Matsumoto in 1994 and in Tokyo

in 1995 [8, 10, 4]. More importantly, by altering the parameters describing the conditions after the attack and the prognosis, the scenario can easily be extended to any event involving a one-time exposure (e.g., chemical agent, bomb explosion, food poisoning). Communicable diseases, radiological releases and events requiring evacuation or quarantine can be captured using additional layers of behavioral and evolutionary complexity.

2. SIGNIFICANCE OF THE SCENARIO

2.1 Sarin and other Nerve Gas Agents

Sarin is a volatile odorless human-made chemical warfare agent classified as a nerve agent [10, 4]. Most nerve agents diffuse because of air currents, sink to lower areas and can penetrate clothing, skin, and mucous membranes in humans. Though Sarin presents only a short-lived threat because of quick evaporation, clothing exposed to Sarin vapor can release Sarin for several minutes after contact.

2.2 Sarin Attacks in Japan

The Aum Shinrikyo cult members initiated Sarin gas release in Matsumoto, Japan on June 27/28, 1994 leading to 7 deaths and injuring over 200. A larger scale attack was executed, less than a year later, on March 20, 1995. The location was a portion of the Tokyo subway system where three train lines intersected and the time was morning rush hour, when the subway was extremely crowded with commuters. Following the attack, all commuters voluntarily evacuated the stations. Emergency Medical Services (EMS) were notified 14 minutes after the event. Police blocked free access to subway stations within an hour. The Japanese Self Defense Forces decontaminated subway stations and trains, and confirmed Sarin as the toxic-agent, three hours after the attack. This 1995 terrorist attack led to 12 fatalities and about 5,500 sickened people [8]. The kinds of questions that analyses can try to address become clear when some of the problems faced in this scenario are considered: (1) overwhelming of communication systems, (2) misclassification and delayed characterization of attack agent, (3) secondary exposure, (4) shortage of hospital resources, (5) lack of mass casualty emergency response plan, (6) absence of centralized coordination, and (7) overwhelming of the medical transportation system.

2.3 Increased Preparedness in Manhattan

The sensational terrorist attack on the Twin Towers of the World Trade Center on September 11, 2001 has made New York city an accessible urban location for analyzing the problems with the emergency response system, warranting well-funded research programs to aid policy development and evaluation. Manhattan, a 20 square mile borough of New York city, is an island in the Hudson River accounting for 1.5 out of the 8 million residents and about 2.9 out of the 8.5 million daytime population. For many reasons, besides the fact that it has become a target of terrorist attacks, Manhattan poses many challenges, serving as an excellent test-bed for verifying assumptions and refining policies about response to large-scale disasters in urban settings. These include: its geographical isolation, tremendous population density (e.g., a day-time population almost double that of the resident population), extensive public transportation system including subways, buses, trains and




Figure 1: Snapshots of the Manhattan model

ferries, its almost vertical structure, its renowned linguistic, ethnic, and socioeconomic diversity, its asymmetric distribution of medical facilities, its proximity to nuclear and toxic-chemical facilities, its ports and airports as an international point of transit and entry, etc. (The model can be seen in *Figure 1*. The color code employed is: person – green(health=1.0), red (health=0.0); hospital/responder – unused (white), inactive (grey), available (blue), critical (pink), full (orange). The streets are black and the subways have the New York subway color codes.)

3. MODELING THE SARIN ATTACK

In this section, we describe the different aspects of our model, the sources of information, the assumptions, the computational approaches and algorithmic issues. Most behavior is probabilistic and most parameters are normalized and initialized uniformly in the range (0, 1).

3.1 Manhattan: Topology and Transportation

We pick the 42nd Street Port Authority Bus Terminal, one block west of Times Square, as the site of Sarin attack. On a typical weekday, approximately 7,200 buses and about 200,000 people use the bus terminal leading to an average flux of over 133 people per minute.

3.1.1 Graph Representation of the Map

The Geographic Information Systems (GIS) street map and the pictorial subway map of Manhattan were obtained from publicly available data sources. The information was converted into a graph with 104,730 nodes (including 167 subway stops) under the following assumptions: (1) Each node represents a location (in real latitude-longitude) where the road curves or where there is a choice of edges to travel on; (2) Each edge represents a straight-line segment of any walkway or a subway; (3) All people and vehicles are constrained to move only along the edges of the graph; (4) The area between streets housing buildings and the area in parks which do not have walkways are deemed unusable for any kind of transportation, even in an emergency; (5) All edges are assumed to be bidirectional. The intersection points were computed assuming that all roads, including flyovers and bridges, intersect all roads that they cross, irrespective of altitude difference. The subway stops were approximated to the nearest node on the graph. The graph is non-planar because of the subway lines, which are mostly underground

in Manhattan. The locations of all major hospitals and some minor hospitals, in all 22 medical facilities, were also approximated to the nearest node on the graph.

3.1.2 Traffic Modeling

Average speed statistics that were available were integrated into a simplistic traffic model. The on-site treatment teams travel at a fixed speed initialized to a random value between 7 and 10 miles per hour. Subways have a fixed speed of 13 miles per hour. Each person has a maximum possible speed initialized to a random value between 6 and 9 miles per hour, consistent with average traffic speed in Midtown Manhattan. To account for congestion, effect of ill-health on mobility and other probabilistic effects, at each time instant, a person travels at an effective speed given by:

```

if(U(0,1) < 1.0-health)
    effective speed = 0.0;
else
    effective speed =
        U(health * maximum speed / 2.0, maximum speed);

```

where $U(0, 1)$ is a real random number generated uniformly in the range $(0, 1)$. No congestion or road width is captured, so there is no enforced maximum number of people at a node or on an edge.

3.2 The People at Port Authority

A “Person” is the most fundamental agent in our multi-agent model, representing the class of individuals exposed to Sarin. However, by-standers and the general population of Manhattan are assumed to play no role (not modeled); same is the case with people and organizations outside the isle of Manhattan.

3.2.1 Person’s Parameters

Based on studies [6, 9] of factors influencing a person’s response to a disaster scenario, the following attributes were chosen to be incorporated into our model: (1) *State*: headed to original destination or to a hospital; (2) *Facts*: current health level (H_i), currently being treated at a hospital or not, current “amount” of medication / treatment, access to a long-distance communication device, probability of the communication device working when the person tries to use it (information update rate); (3) *Knowledge*: location and current capacities of known hospitals and on-site treatment units, time of last-update of this information, tables of the *LRTA** estimates for the known nodes, list of 100 most recently visited nodes; (4) *Personality*: degree of worry (W_i), level of obedience (O_i), perceived level of distress ($D = W_i \times (1 - H_i)$). The obedience parameter O_i captures the instruction-abiding trait of a person, and affects the decision to head to a hospital. The worry parameter W_i represents the innate level of irrationality in the agent’s behavior, and affects the following decisions: when to go to a hospital, when to get information from neighbors or via cell phone, or how to select the hospital.

3.2.2 Rules of Behavior

The person’s initial goal is to reach the original destination (e.g., home or place of work) from the initial location (the Port Authority Bus Terminal). However, after exposure to Sarin, his/her health begins to deteriorate. At a certain health-level decided by environmental and person-

ality factors, the person changes the destination state to a hospital:

```

if(U(0,1) < Obedience) {
    if (health < unsafe health level)
        Head to a hospital
    }
else if (U(0,1) < distress level)
    Head to a hospital
}

```

where the unsafe health level is the suggested health level when a person should head to a hospital.

Initially, each person agent knows only a random number of hospitals and their absolute positions in the map (latitude and longitude), but this knowledge gets updated during the evolution of a simulation using the different communication channels (described in *Section 3.5*):

```

if (heading to a hospital && U(0,1) < distress level) {
    if (U(0,1) < information update rate)
        Get current hospital information via phone/radio
    else
        Talk to neighbors
}

```

The choice of hospital is then made based on the list of hospitals and on-site treatment facilities known, their current capacities, and personality and environmental factors:

```

if(U(0,1) < distress level) {
    Find nearest hospital
} else {
    Find nearest hospital in available mode
}

```

After being treated and cured at a medical facility, the person resumes moving towards his/her original destination.

3.2.3 LRTA* with Ignore-List for Route Finding

The Learning Real-Time (*LRTA**) algorithm, proposed by Korf in 1990 [5], interleaves planning and execution in an on-line decision-making setting. In our model, the person-agent is modeled as maintaining an “ignore-list” of the last 100 nodes he/she visited, and employs the following *modified LRTA** algorithm:

1. *Default*: If all neighbors of the current node i are in the ignore list, pick one randomly.
2. Else:
 - (a) *Look-Ahead*: Calculate $f(j) = k(i, j) + h(j)$ for each neighbor j of the current node i that is not in the ignore-list. Here, $h(j)$ is the agent’s current estimate of the minimal time-cost required to reach the goal node from j , and $k(i, j)$ is the link time-cost from i to j , which depends on the type of the link (road or subway) and its effective speed (subway or person speed).
 - (b) *Update*: Update the estimate of node i as follows:

$$h(i) = \max\{h(i), \min_{j \in Next(i)} f(j)\}$$

- (c) *Action Selection*: Move towards the neighbor j that has the minimum $f(j)$ value.

As the planning time for each action executed by the agent is bounded (constant time), the *LRTA** algorithm is known to be usable as a control policy for autonomous agents, even in an unknown or non-stationary environment. However, the rational *LRTA** algorithm was inappropriate in its direct form for modeling persons trying to find the route to their original destination or hospital in an atmosphere of tension and panic. Thus, the *ignore-list* was introduced to capture a common aspect of panic behavior: people seldom return to a previously visited node when an unexplored node is available. In other words, the only case when a person uses old learnt information is when they revisit a node they visited over a hundred nodes ago. The algorithmic characteristics of this “ignore-list” heuristic are being investigated separately.

3.3 The Medical Facilities in Manhattan

The hospital agent is a stationary agent that is an abstraction of any medical facility that can play a role at the time of a catastrophe. Twenty two major and minor hospitals have been included, and the number of hospital beds was used as an indicator of the capacity (“resources”) of the hospital.

3.3.1 Hospital’s Parameters

The attributes of a hospital that are included in our model are: (1) *State*: available, critical or full; (2) *Facts*: resource level (representing both recoverable resources like doctors, nurses and beds, and irrecoverable resources like drugs and saline), reliability of communication device (information update rate); (3) *Knowledge*: locations and current capacities of known hospitals; (4) *Triage Behavior*: health-levels below which a person is considered critical, non-critical or dischargeable.

3.3.2 Rules of Behavior

As described in our Brazilian scenario model [7], the hospital operates in three modes: “available”, “critical” and “full”. When a hospital’s resource level drops below the *low* resource level ($\frac{1}{3}^{rd}$ of initial resources), its mode changes from available to critical. When a hospital’s resource level drops below the *very low* resource level ($\frac{1}{10}^{th}$ of initial resources), its mode changes from critical to full. The hospital mode directly influences the key decisions: whom to turn away, whom to treat and how much resources to allocate to a person requiring treatment. The medical parlance for this process is “triage”, and research is actively being conducted to evaluate different triage policies appropriate to different scenarios (for example, see the Simple Triage and Rapid Treatment system [10]). The hospital’s behavior at each time step is described by the following rules:

```
Treat all admitted patients
for all persons inside the hospital{
  if (health >= dischargeable health level)
    Discharge person
  else if(person is waiting for admission) {
    if(hospital is in available mode)
      Admit and treat the person
    else if(hospital is in critical mode &&
      health < critical health level)
      Admit and treat the person
  }
  if (person is waiting &&
    health < critical health level)
    Add to critical list
```

```
  if (person is admitted &&
    health > non-critical health level)
    Add to non-critical list
}
Discharge non-critical patients, admit critically ill
```

3.4 On-Site Treatment Units

On-site treatment is provided by Major Emergency Response Vehicles (MERVs) which set up their units close to the site of action. The HazMat Team consists of experts trained in handling hazardous materials, who rescue people from the contaminated zone, collect samples for testing, and eventually decontaminate the area. In our model, we group HazMat and MERVs into one unit – “on-site treatment providers”. These small mobile hospitals are initially inactive and stationary at their hospital of affiliation. When notified of the attack, they move towards the catastrophe site. Their properties include: (1) *Facts*: starting location, time of dispatch; (2) *Knowledge*: locations and current capacities of known hospitals; tables of the *LRTA** estimates for the known nodes, list of 100 most recently visited nodes; (3) *Behavior*: exactly the same as a hospital in “critical” mode;

The model for which the statistics are reported in this paper has 5 on-site treatment providers. In a real situation, the first responders to the emergency include the Police and Fire department personnel. Ambulances arrive at the scene and transport sick people to the hospitals. No ambulance-like services are currently part of the model. The role of the police in cordoning the area and crowd management is implicit in that on-lookers and by-standers do not complicate the disaster management process in our model.

3.5 Communication Channels

In the model analyzed in this paper, only the information about the hospital and on-site treatment provider locations and capacities are communicated dynamically. The channel of communication used for on-site treatment provider activation is not modeled; only the time of availability of the information is controlled. The communication channels available are: *one-to-one* between persons and any of the other three classes of agents adjacent to them, *one-to-many* from the hospital to all persons within its premises, and *many-to-many* from the hospitals to all other hospitals, persons and on-site treatment units with access to a public telephone, radio or a mobile communication device. The role of media, internet, misinformation and rumors are not modeled.

3.6 Sarin Gas Exposure

3.6.1 Time-course of Deterioration and Recovery

The time-course variation of the health level (with and without treatment) after the exposure is modeled using a 3-step probabilistic function depending on the person’s current health level.

```
if (U(0,1) < health) {
  health = health
  + U(0, treatment + maximum untreated recovery);
else
  worsening = (health > dangerous health level)?
    maximum worsening:
    ((health > critical health level)?
      maximum dangerous worsening:
      maximum critical worsening)
  health = health - U(0, (1 - treatment)*worsening);
```

Table 1: Exposure level and health level ranges

Exposure level	Health range	People Exposed
High (lethal injuries)	(0.0, 0.2]	5%
Intermediate (severe injuries)	(0.2, 0.5]	25%
Low (light injuries)	(0.5, 0.8]	35%
No symptoms	(0.8, 1.0)	35%

The exact values used are dangerous health level = 0.5, critical health level = 0.2, maximum worsening = 1.38×10^{-4} per minute, maximum dangerous worsening = 4.16×10^{-4} per minute and maximum critical worsening = 6.95×10^{-4} per minute.

3.6.2 Level of Exposure

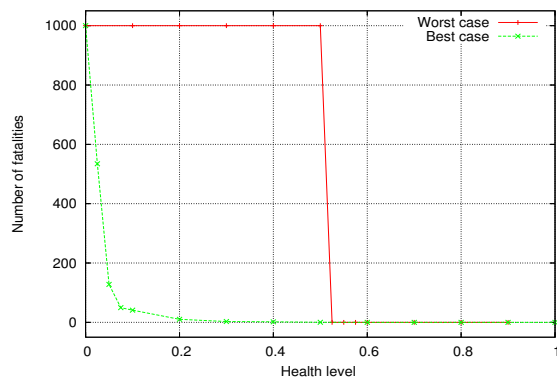
Based on diffusion effects, air-currents, number of people, temperature, time of day, rate of breathing and amount of time exposed to Sarin, the amount of Sarin inhaled by a person (“acquired dose”) at a certain distance from the source can be estimated. Based on this dosage, a certain health response results (based on “dose-response curves” in toxicology). Unfortunately, it is impossible to estimate the nature, intensity and location of an attack (even within the Port Authority Bus Terminal). More importantly, there is no clear-cut data on the rate of health degradation after exposure to a certain dosage. This is significant, as the ultimate aim of the modeling is to see how the time taken by the on-site responder units to initiate treatment compares with the time taken by the Sarin poisoning to result in death. Reasonable estimates for the rate of health deterioration were arrived at in consultation with toxicologists in the CCPR team and based on related literature [10, 4]. Table 1 shows the four main classes of exposure that have been modeled, the corresponding ranges of initialization for the health level and the percentage of people initialized to that category. These values reflect our attempt to capture the general situation of previously documented events[8], where only a small fraction of the affected population suffered fatal injuries. One key assumption in our model is that there is no secondary exposure, i.e., on-site treatment units and hospital staff are not affected by treating Sarin-exposed patients.

3.6.3 Chances of Survival

The actual survival chances under optimistic and pessimistic conditions that result from the assumptions of our model are depicted in *Figure 2*. People with fatal and severe injuries can survive if they are treated on-site or if they are transported to a nearby hospital. People with light injuries and those showing no symptoms will always recover eventually, but in this case, the damage to organs and the time to recover are the correct metrics of effectiveness of the emergency response. However, in this paper, we focus only on the number of deaths. As the survival-chances curve shows, only people with health less than 0.5 can ever die. However, all persons factor in, as they decide how information percolates and how resources are distributed.

4. ANALYSIS OF SIMULATIONS

Since no well-defined approaches exist for choosing the correct level of abstraction and identifying the essential parameters for modeling a scenario, a significant portion of agent-based modeling remains an art more than a science.

**Figure 2: Sarin: Treatment and Survival Chances**

The assumptions used in our model, made in consultation with experts from the CCPR team and based on related literature, were often made for want of accurate data or for simplification of the analysis. It is reiterated that the simulations cannot by themselves serve as factual outcomes, and so, emergency response planners are expected to integrate scientific expertise, field exercises and historical data with these simulation results to make sound decisions in real scenarios.

The model has been implemented in the Java version of RePast 3.1[2], a popular and versatile toolkit for multi-agent modeling. In the results described below, the following additional assumptions were made: (1) The simulation is performed only for the first 3000 minutes (= 2 days and 2 hours). The assumption is that people who survive the first two days are not likely to die. Further, by this time resources from the outside the island of Manhattan will become available and the scenario is beyond the scope of our current model; (2) Neither an on-site responder nor a hospital can help a person if the person does not ask for treatment (“head to a hospital” mode); (3) None of the behavior parameters change during a simulation, as learning behavior is supported only for the route finding algorithm. Unless stated otherwise, all plots involve 1,000 people, 22 hospitals, and 5 on-site responder teams. Every point that is plotted is the average of 10 independent runs. All plots without responders start at a slightly different initial state (with identical stochastic properties).

4.1 People Behavior

4.1.1 Unsafe Health Level

A critical disaster management question is: When should a person experiencing symptoms go to a hospital? Consider the scenario when there are no on-site treatment units. In *Figure 3*, the influence of the health-level at which a person decides to go to a hospital (called “unsafe health level”) on the number of deaths is visualized. This plot suggests that person should decide to go to a hospital when his or her health approaches 0.2.

This unexpectedly low optimum value reflects a skewed health scale and can be explained thus. From *Figure 2* we observe that if the health level is greater than 0.1, almost 95% of the people will recover fully with treatment, while if the health level is greater than 0.5, 100% of them will recover even without any treatment. When the unsafe health

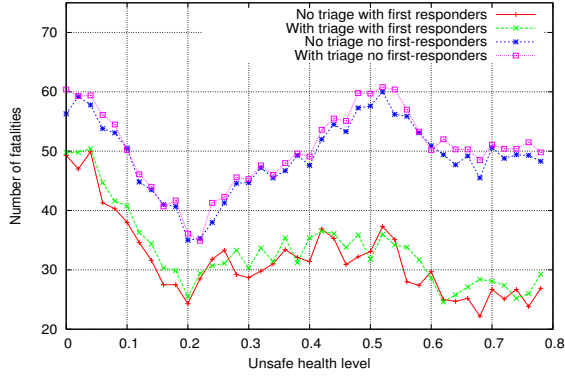


Figure 3: Persons heading to a hospital with and without on-site treatment units (number of on-site responders = 5, on-site responder’s dischargeable health level = 0.5, hospital’s dischargeable health level = 0.8, responder alert time = 15 minutes).

level is too low (< 0.2), people have been instructed to wait so much that their condition turns fatal. The second factor affecting the optimum value for heading to a hospital is the distribution of people across the different classes of injuries. As seen in *Table 1*, a cut-off of 0.2 ensures that only the people who experienced lethal injuries (50/1000) go to a hospital. The moment this cut-off is increased, to say 0.5, crowding effects hamper emergency response as another 250 severely injured persons also rush to the hospitals. This situation is exacerbated by the fact that health level governs mobility, and hence healthier people are expected to reach a hospital earlier than sicker people. Thus, when unsafe health level is high (> 0.2), people who do not require much emergency treatment end up consuming a share of the available resources, which would have been better spent on the sicker people already at the hospital or on persons who are still on their way to the hospital. Clearly, the presence of ambulances would alter the situation as the lethally injured persons would actually move faster than persons of all other classes. The drop in death rate after 0.6 can be attributed to the fact that people with health level greater than 0.6 would have recovered by themselves (see *Fig. 2*) on the way to the hospital, and hence may have not applied any pressure on the hospital resources.

The number of deaths due to crowding is dramatically mitigated if there are on-site treatment units, as seen in *Figure 3*. It is to be recalled that from the point of view of a person, an on-site treatment unit is equivalent to a hospital in “critical” mode. The number of deaths due to people heading to hospitals earlier than necessary is less, as most of these very sick people are now treated on-site, and hence, are no longer dependent on the resources of the hospitals. When a person’s health level is greater than the unsafe health level, in addition to not heading to a hospital, the person refuses treatment even from an on-site treatment provider. Though this assumption is unrealistic when the person’s health is less than 0.2 or so, it is plotted for completeness.

4.1.2 Worry and Obedience

Two significant personality parameters that affect disaster-time behavior of a person are the innate degree of worry and




Figure 4: Effect of people’s obedience and worry levels (hospital’s dischargeable health level = 0.8).

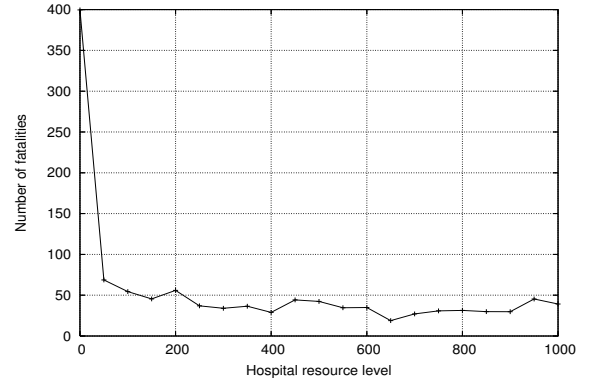


Figure 5: The effect of having more resources

obedience (see *Sec. 3.2.2*). These population parameters can be controlled by education, awareness and training before an event, and also by employing law enforcement officers during the emergency response. Obedient persons do not head to a hospital when their health level is above what is considered unsafe, while disobedient persons will go based on their perceived level of distress. In order to understand their influence on the global system behavior, a set of simulations were performed by varying both O_i and W_i in the range $[0, 1]$ and assuming that on-site responders are not active. *Figure 4* shows the results of their mutual interaction. By our definition of obedience and worry, disobedient worrying persons will head to the nearest hospital too early, thus crowding the most critical resource. At the other extreme, obedient people who are not worried choose to go to a hospital only when they are really sick, and also distribute themselves between the different hospitals; only when they become critically ill do they go to the nearest hospital irrespective of its mode. Disobedient people who are not worried do not worsen the situation because they will still get hospital information and choose to go to one, only when necessary (based on level of ill-health).

4.2 Hospital Behavior

4.2.1 Resource Requirements

The meaning of the “resource” parameter is clarified in *Figure 5*. The thought experiment that led to this plot was:




Figure 6: Hospital’s patient-discharge behavior without on-site treatment (Person’s unsafe health level = 0.2).

when there is only one hospital, and the Sarin attack occurs immediately adjacent to it, how much resources are necessary to save the 1000 affected people? As the plot shows, if the hospital has resources > 100.0 , then no more than 50 deaths can result. A resource level > 200.0 can bring the number down between 40 or 20. The number of deaths is never zero because the personality parameters make different people choose to head to the hospital at different times.

4.2.2 Optimal Dischargeable Health Level

The hospital’s decision to discharge a patient is dictated by its estimate of whether the patient can recover using just medication, as opposed to requiring continuous monitoring. In our model, the hospital discharges persons whose health level is greater than “dischargeable health level”. In *Figure 6*, the relationship of this decision with the number of deaths is plotted, and is seen to follow the same pattern as the “unsafe health level”. When the dischargeable health level is too low, the person dies after being discharged prematurely. When it is too high, the person is given more medical attention than necessary and effectively decreases the chances of survival of sicker persons.

It is not immediately clear why the death-rate drops when the dischargeable health level is greater than 0.6. One possible explanation is that a person so discharged always recovers fully, whereas a fraction of the people discharged earlier return for treatment, possibly to a different hospital. The peak near 0.0 of 50 deaths is less than the peak near 0.6 of 65 deaths. This is because the hospital in reality is not entirely refusing treatment to persons with health level greater than dischargeable health level. The person is given *some* treatment, then discharged and then readmitted until the person’s health becomes greater than the unsafe health level, at which point he/she “accepts” the hospital’s decision to discharge him/her and resumes moving towards his/her original destination. Also, unpredictable behaviors can result when the linear ordering of the parameters ($0 < \text{critical health level} < \text{non-critical health level} < \text{dischargeable health level} < 1$) is violated.

The behaviors with and without triage not being very different may be related to the fact that hospitals broadcast their mode irrespective of whether they are enforcing triage policies or not. Persons use this information to choose the




Figure 7: Number of on-site responders versus their alert time (on-site responder’s dischargeable health level = 0.5, hospital’s dischargeable health level = 0.8, person’s unsafe health level = 0.4).

hospital. Since we are counting only the number of deaths and since the very sick people go to the nearest hospital irrespective of triage enforcement, only the difference in the behavior of the hospital affects the result. However, in the critical mode, the hospital admits all persons with health level less than the critical health level ($= 0.25$). Thus the differences are minimal when the triage is enforced and the hospital is in the critical or available mode. The difference would have been noticeable had the hospitals been smaller or had the number of people been more; then the hospitals would have moved to “full” mode refusing admission to even the critically ill.

4.3 Role of On-Site Treatment Providers

The role of the on-site treatment providers is patent in the parameter surface (*Figure 7*) of their alert time versus the number of fatalities. As expected, the plot also shows a near-linear dependence of the system on the alert time. However, beyond the 10 responders that seem to be required, the effect on the improvement in the number of survivors is less evident. Clearly, the bound on the number of dying people that can actually be saved causes this surface flattening.

4.4 Significance of Communication

4.4.1 Getting Current Hospital Information

We modeled the scenario where every person has a communication device, and then controlled the rate of information update (which can capture the difficulty of access to a cell-phone, taxi-phone or public phone booth, the congested nature of the communication network, the lack of response from the callee, etc.). The impact of this parameter on the number of fatalities is plotted in *Figure 8*. As observed in the Brazilian scenario analysis[7] also, the death rate declines when more people have complete hospital information. When everybody has access to the current information about all hospitals, healthier people, who would have survived the commute to a farther hospital, consume the resources of the nearby hospitals which they quickly reach. Critically ill people, whose average speed is much lower, are effectively left with a choice only between critical or full proximal hospitals and available distant hospitals – both of which turn out to be fatal.




Figure 8: Person’s ability to communicate (person’s unsafe health level = 0.2).

4.4.2 Activating the On-Site Responders

The success of the on-site treatment responders is dependent on how quickly they get alerted, as shown in Figure 7. As a result of our parameter choice, we see that the net number of fatalities is stable (~ 25), as long as the on-site responders arrive within 50 minutes. The fluctuations could be due to the fact that the persons are themselves moving and need to be able to locate the on-site responder.

5. DISCUSSION

Several important emergency response issues, such as when to head to a hospital, when to discharge a person, number of on-site treatment units necessary, the importance of public awareness and law enforcement, the role of responder size and activation time, and the diffusion of information about hospitals and capacities, were amenable to analysis by repeated simulation. ABM shows tremendous potential as a simulation-based tool for aiding disaster management policy refinement and evaluation, and also as a simulator for training and designing field exercises.

The “Sarin in Manhattan” model in itself can be extended by addressing the assumptions discussed earlier. On the computational side, better knowledge and belief-state representation are necessary to simplify and generalize the communication mechanisms. Further, this will lead to simpler encoding of learning behavior; thus all parameters, including personality states, should be able to evolve with experience. We modified the simple *LRTA** algorithm to take into account the memory of recently visited nodes to approximate real human panic behavior. This model needs to be refined, and more personality and learnt parameters need to be factored in. Another aspect that is missing in our model is the information about routes and location of subway stops. After modeling traffic congestion, the role of a centralized navigation system [12] in managing disaster-time traffic and routing also warrants investigation. To improve the ultimate utility of the tool, we need to devise a uniform way of describing different catastrophic scenarios, with the ability to validate over well-documented real instances. Further, a conventional AUML-based description of agent behavior needs to be the input for the system. Some of the specific scenarios we hope to model in the near future include food-poisoning, moving radioactive cloud, communicable diseases, natural disasters leading to resource damage

in addition to disease, and events requiring evacuation or quarantine. On the theoretical side, we would like to automate the process of policy evaluation and comparison, and optimal parameter value estimation. We are also investigating representations of plans so that multi-objective optimization via evolutionary algorithms can be used to design new emergency response strategies. To address cultural and racial differences in response to catastrophes, game-theoretic behavior modeling and analysis is being surveyed [1].

6. ADDITIONAL AUTHORS

Additional authors: Lewis Nelson (NYU School of Medicine, email: lnelson@pol.net), Dianne Rekow (NYU College of Dentistry, email: edr1@nyu.edu), Marc Triola (NYU School of Medicine, email: marc.triola@med.nyu.edu), Alan Shapiro (NYU School of Medicine, email: alan.shapiro@med.nyu.edu), Clare Coleman (NYU CCPR, email: colemc01@med.nyu.edu), Ofer Gill (NYU, email: gill@cs.nyu.edu) and Raoul-Sam Daruwala (NYU, email: raoul@cs.nyu.edu).

7. REFERENCES

- [1] G. Boella and L. W. N. van der Torre. Enforceable social laws. In *AAMAS 2005*, pages 682–689. ACM, 2005.
- [2] N. Collier, T. Howe, R. Najlis, M. North, and J. Vos. Recursive porous agent simulation toolkit, 2005.
- [3] E. A. der Heide. The importance of evidence-based disaster planning. *Annals of Emergency Medicine*, 47(1):34–49, 2006.
- [4] C. L. Ernest. Clinical manifestations of sarin nerve gas exposure. *JAMA*, 290(5):659–662, 2003.
- [5] R. Korf. Real-time heuristic search. *Artificial Intelligence*, 42:189–211, 1990.
- [6] R. Lasker. Redefining readiness: Terrorism planning through the eyes of the public, 2004.
- [7] V. Mysore, O. Gill, R.-S. Daruwala, M. Antoniotti, V. Saraswat, and B. Mishra. Multi-agent modeling and analysis of the brazilian food-poisoning scenario. In *The Agent Conference*, 2005.
- [8] T. Okumura, K. Suzuki, A. Fukuda, A. Kohama, N. Takasu, S. Ishimatsu, and S. Hinohara. The tokyo subway sarin attack: disaster management, part 1: Community emergency response, part 2: Hospital response. *Academic Emergency Medicine*, 5(6):613–24, Jun 1998.
- [9] PNAS. *Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling*, volume 99(3). May 2002.
- [10] Y. Tokuda, M. Kikuchi, O. Takahashi, and G. H. Stein. Prehospital management of sarin nerve gas terrorism in urban settings: 10 years of progress after the tokyo subway sarin attack. *Resuscitation*, 68:193–202, 2006.
- [11] I. van Heerden and B. A. Hurricane pam exercise - july 2004, adcirc storm surge animations and related information, 2004.
- [12] T. Yamashita, K. Izumi, K. Kurumatani, and H. Nakashima. Smooth traffic flow with a cooperative car navigation system. In *AAMAS 2005*, pages 478–485. ACM, 2005.