Modeling Human Behavior in the Aftermath of a Hypothetical Improvised Nuclear Detonation

Nidhi Parikh, Samarth Swarup, Paula E. Stretz, Caitlin M. Rivers, Bryan L. Lewis, Madhav V. Marathe, Stephen G. Eubank, Christopher L. Barrett, Kristian Lum, Youngyun Chungbaek
Network Dynamics and Simulation Science Lab,
Virginia Bioinformatics Institute,
Virginia Tech,
Blacksburg, VA, USA
{nidhip, swarup, pstretz, cmrivers, blewis, mmarathe, seubank, cbarrett, klum, ychungba}@vbi.vt.edu

ABSTRACT

In this paper we describe a multiagent simulation model of human behavior in the aftermath of a hypothetical, large-scale, human-initiated crisis in the center of Washington D.C.

Prior studies of this scenario have focused on modeling the physical effects of the attack, such as thermal and blast effects, prompt radiation, and fallout. Casualty and mortality estimates have been obtained by assuming a spatially static human population, ignoring human behavioral response to the event.

We build a simulation of a behaving human population and its interaction with various interdependent infrastructures, to try to understand how human response to such an event would change outcomes, and also how modeling this response would enable us to develop new perspectives on planning for this event.

Here we present details of the simulation, focusing on the agent design and multiagent interaction, and present initial results on how rapid restoration of communication could alter behavior beneficially.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences

General Terms

Algorithms, Experimentation

Keywords

Social Simulation; Behavior Modeling; Disaster Modeling

1. INTRODUCTION

Human population is subjected to risk due to various types of life-threatening events like fire, flood, earthquake, accident in a nuclear plant or terrorist attack. Though various offices and buildings conduct fire-drill, tornado-drill and other mock training to prepare their employees or people who reside in those buildings, it is nearly impossible to carry out such mock scenario training when a large population (e.g., that of a city) is involved in a life-threatening situation. Planning and responding to such major disasters is a challenging task as it not only requires taking into account the physical impact but also knowing geographical distribution of the population, and accounting for their behavior. Agent-based simulation approach involves modeling actions and interactions of autonomous agents to forecast emergent collective behavior and hence could help aligning response policy with survivors' behavior.

The primary goal in any emergency situation is to save lives and minimize the extent of injuries, while reducing the damage to infrastructure is a secondary goal. However in reality human behavior and various infrastructural systems are coupled – what people do creates a load on infrastructural systems, and availability and damage to the infrastructure restrict or alter the behavior of individuals, as illustrated in fig. 1, for example.

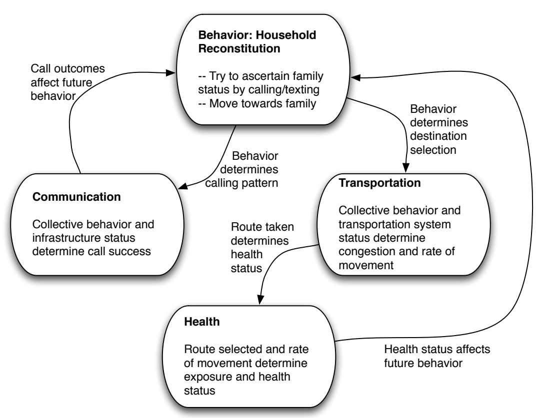

Figure 1: Human behavior interacts with infrastructure in feedback loops.

This work is part of a project that models human interactions and behavior (including its impact on health) coupled with multiple infrastructural systems including the power network, the communication (cell phone) network, and the transportation network (including road, metro and bus networks) in the aftermath of the detonation of a nuclear device in Washington DC. This paper particularly focuses on

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

agent behavior and their interactions with each other and with these infrastructural systems and its impact on health.

2. RELATED WORK

Earlier work on the effect of the detonation of a nuclear device [5, 22] or dirty bomb [6] has been focused on evaluating evacuation vs. shelter in place policies. They model the physical impact of the blast in terms of thermal, radiation, and fallout effects and use static geographic distribution of the population during daytime or at night time (from Landscan data, e.g.) to calculate effects on human life. However, they ignore human behavioral response to the event, like family members looking for each other, survivors carrying out search and rescue, and so on.

Multiagent models allow each individual to be modeled as an autonomous agent capable of perceiving informational and environmental cues and interacting with other agents and the environment accordingly. As a result, agent-based methods have been used for evacuation simulations [12, 14, 21]. Simplistic models focus on physical interactions between individuals [12]. Somewhat more detailed simulations model agents with psychological models, though they haven't included infrastructural aspects [16, 19]. Tsai et al. [21] and Pan et al. [14] model individual movement towards building exits and behavioral aspects like family influence and following the group leader in simulations of emergency evacuation of indoor spaces. Our goal here is not only to model evacuation related movement and behavior but also to model other behavioral characteristics of humans in crisis (like helping others, sheltering in place, etc.) to understand the interactions between natural human behavioral instincts and external interventions.

There have been multiple surveys and analyses, both prospective and retrospective, of human behavior in disasters. They provide insight into different kinds of behavior:

- *Sheltering, seeking family members, communication*: Various surveys [9, 10] indicate that most people would try to leave the area if not asked to shelter in place in case of a dirty bomb. The main reason for individuals to leave is concern for people dependent upon them and other family members [9, 10, 11]. However, people would stay in-place if they are able to communicate with their family members [10].
- *Evacuating only after finding family members*: If an emergency happens during day time, members of a family are likely to be scattered across the region (for daily activities like work, school, etc.). In such cases family members try to gather children [13] and each other and evacuate as a single unit [7].
- *Delay in evacuation*: It is also observed that in an emergency without warning (such as terrorist attack), there is some delay between the time at which the initial cues occur that an emergency is taking place and the time people start evacuating based on their perception of risk, which is based on environmental cues (i.e., smoke, debris), behavior of others, and past experience [18].
- *Aiding and assisting, seeking healthcare*: Contrary to the assumption that trained emergency personnel carry out field search and rescue, studies show that most initial search and rescue is carried out by survivors [2, 17]. Also survivors and most casualties are more likely to go to the nearest hospital.

We use these findings to build the behavior model for our agents, as detailed in section 5.

3. SCENARIO

We build on the work of Buddemeier et al. [5] and Wein et al. [22] to construct the scenario for the simulation. The hypothetical detonation occurs on a weekday morning on the corner of 16^{th} and K Street NW in Washington DC. The fallout cloud spreads mainly eastward and east-by-northeastward.

We simulate the population inside a region we call the detailed study area (DSA; fig. 2), which is the area defined by the .01 Gy fallout contour at 60 minutes joined with the thermal radiation contour at 2.1 cal/cm^2 bounded by the boundary of the counties neighboring the District of Columbia.

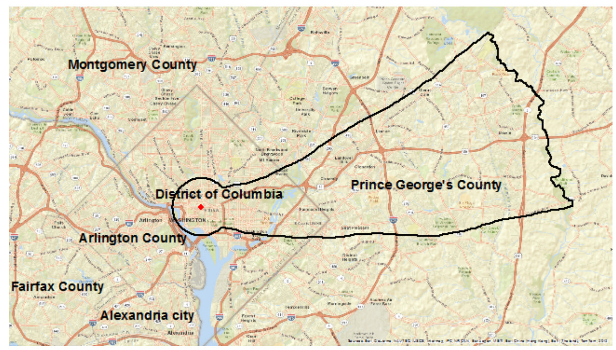

Figure 2: The detailed study area (DSA).

The blast causes significant disruptions in the power system and cell phone system, and significant damage to buildings and roads. The full simulation uses detailed data about each of these infrastructures to create models of phone call and text message capacity, altered movement patterns due to road damage, injuries due to rubble and debris, and levels of radiation protection in damaged buildings. We omit most of these details from the present paper due to lack of space[1], in order to focus on agent design, agent interactions, and evaluation.

4. POPULATION MODELING

Our simulation uses a *synthetic population* [4, 3] of the Washington DC metro area (which includes surrounding counties), which we have extended to include transients.

The scenario affects all people present in DSA at the time of the event, which includes residents, tourists, business travelers, and college students. Their health and behavior in the aftermath of the event depend upon where they are located when the event happens. Hence a detailed synthetic population has been modeled in a way that in addition to giving information about demographics, also includes information about the daily routine of each individual. The process for creating the synthetic population is outlined below:

[1]More details are available at `http://ndssl.vbi.vt.edu/projects/disaster-resilience/`

4.1 Base Population (Residents)

Generating the Population: Demographic distributions and sample household information from the American Community Survey (ACS) are used to create a disaggregated population, which consists of a set of synthetic households and a set of synthetic individuals. This is done by using an algorithm known as iterative proportional fitting to generate a joint distribution, which is then sampled [4]. The generated synthetic population matches marginal demographic distributions from the ACS at the block group level, while preserving anonymity of individuals.

Locating Households: Each synthetic household is assigned a housing location along a street using housing unit distributions from the ACS and street data from Navteq.

Assigning Activities: Each individual is assigned a set of activities to perform during a day, along with the time. The National Household Travel Survey (NHTS) and the National Center for Education Statistics (NCES) are used to create activity templates. Each synthetic household is matched to a survey a household based on its demographics and individuals in synthetic household are assigned the corresponding activities.

Locating Activities: An appropriate location (essentially a building) is chosen for each activity of each individual using a gravity model and Dun & Bradstreet location data.

4.2 Transient Population

The method used to create the transient population [15] follows the same methodology as used for creating the base population, but using different data sources. Destination DC[2] provides demographic information about leisure and business travelers in Washington DC. These demographic data are used to create synthetic population of transients. Destination DC estimates that there are approximately 50000 transients in Washington DC on any given day. The transient population agents are divided into groups called parties, e.g., a family of tourists traveling together. Each party is placed in a hotel which serves as their home location for the purpose of visit. All individuals in a party are assumed to travel together and hence assigned the same activities. Each activity is represented by the type of activity (i.e., staying in hotel, tourism, going to restaurants, work in case of business travelers), start time, duration and location. Various activity locations have been identified from Dun & Bradstreet data based on SIC (Standard Industrial Classification) codes. Activity assignment is calibrated by matching visit counts at Smithsonian Institution locations, which are the largest draw for tourists.

4.3 Dorm Student Population

A synthetic population of college students living in dormitories is created separately for major colleges in the DSA. Data about the number of dorm students in each college and college boundary are obtained from CityTownInfo[3] and the District of Columbia public access online Data Catalog[4] respectively. For simplicity, students are assigned only two types of activities, staying in the dorm and school activities located at any of the locations within their college campus.

[2]http://washington.org

[3]http://www.citytowninfo.com/

[4]http://data.dc.gov

Table 1: Datasets used for population generation.

Used for	Data source
Base US population	American Community Survey National Center for Education Stat. National Household Travel Survey Navteq Dun & Bradstreet
Transient population (additional)	Destination DC Smithsonian visit counts
Dorm students (additional)	CityTownInfo District of Columbia public access online Data Catalog

All the data sets used are summarized in table 1. The total size of the synthetic population for the Washington DC metro area is over 4 million. The total number of agents in the simulation is 730,833, which is the subset of individuals within the DSA at the time of the event, and the total number of locations within the DSA is 146,337.

5. AGENT DESIGN

In addition to the demographic variables described above, agents are defined by a number of state variables, which are of three kinds: the agent's knowledge of his family members' health states, the agent's current "behavior", and the agent's own health state. We describe each of these in turn.

Knowledge of family status: For each family member, an agent keeps track of whether their health status is unknown, known to be healthy, or known to be injured. These variables get updated if the agent is able to establish communication with a family member, through a phone call or text message, or if they encounter each other in person.

Follow the leader behavior: When agents encounter their family members, they are grouped with them, so that they move on as a unit. This is done by choosing one of the agents as the group leader who then makes decisions for the entire group. This sort of behavior in emergency situations is well-documented in the literature [7, 13]. Similarly, if an agent is assisting another agent, even if the second agent is a non-family member, they will be grouped together, with the first agent as the group leader in order to indicate that they are traveling together. It is known in the literature on human behavior in emergencies that most of the initial search & rescue work is carried out by civilians [2]. Agents who are assisted or rescued by other agents in the simulation are ones who are too injured to move by themselves, or are small children. This behavior is explained in more detail in section 5.2.6.

5.1 Agent Behavior

The behavior model for the agents is based on the decentralized Semi-Markov Decision Process (Dec-SMDP) formalism [8, e.g.]. We use the framework of *options*, where each option is a policy together with initiation and termination conditions [20]. Options define high-level behaviors. In our model, agents can choose between six options: household reconstitution, evacuation, shelter-seeking, healthcare-seeking, panic, and aid & assist.

A policy is specified as a short program for selecting action. Actions consist of attempting to move toward some destination, or attempting to establish communication with someone. The destination chosen depends on the option being executed, as does the calling behavior. For example, in the healthcare-seeking option, agents may try to move towards hospitals and call 911, whereas in the household reconstitution option, they may try to call their family members and move towards them. All the options are presented in detail below.

At each time step, agent behaviors are updated by first checking the termination condition for the current option, then choosing a new option if the current one is terminated, and then choosing an action based on the current option. Formally, the option selection mechanism is just required to be a probability distribution. However, to make it more human-interpretable and also as a means of embedding prior knowledge about which behaviors are reasonable in which circumstances, we represent the option selection mechanism as a decision tree, where each leaf contains a probability distribution over the options considered available in those conditions. This decision tree is shown in fig. 3.

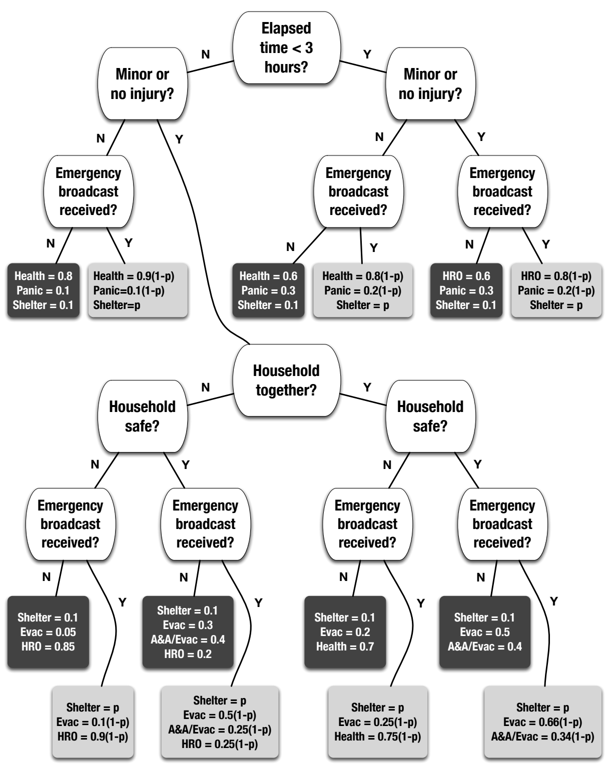

Figure 3: Decision tree for behavior option selection.

The set of options is chosen based on the literature, as described in section 2. In addition, the option selection decision tree in fig. 3 incorporates the following assumptions.

People are more likely to panic soon after the event as compared to later. We assume that after 3 hours, levels of panic begin to fall naturally, but this can happen faster if people obtain information about the event, e.g., through emergency broadcasts.

As mentioned in section 2, there tends to be a delay between the time at which initial cues about the emergency occur and the time that people start evacuating. We assume that this delay is also 3 hours, and that people evacuate only if they are healthy enough to do so, and have managed to locate their family members or ascertain that they are safe. Aid & assist behavior is delayed for the same reason and is only chosen by agents close to ground zero. Agents further away choose to evacuate instead.

It is assumed that authorities start sending out emergency broadcasts advising people that a nuclear detonation has occurred and that they should seek shelter. This has repeatedly been suggested as the best course of action during such an event [10, 11, 22, 5]. In our option selection tree, we make the probability that an agent shelters upon receiving an emergency broadcast a parameter, p.

For each high level behavior, an agent performs an action (i.e. call, move, or both). Individual movement is taken care of by a transportation model which takes into account road, bus and metro network and damage to them. Calls attempted are taken care of by wireless communication model that takes into account phone availability, reception, and available bandwidth. Detailed description of these systems are out of the scope the current paper.

It is hard to obtain exact probabilities for each behavior option and action in all circumstances from the literature. The values used for the simulation are shown in the figures, however they could be changed if necessary.

5.2 Behavior Options

We now describe each of the behavioral options, including the policy and the termination conditions. The initiation conditions are as shown in the option selection decision tree in fig. 3.

5.2.1 Household Reconstitution Option (HRO)

Household reconstitution (seeking family members or information about them) is the most natural human behavior in emergency situations.

Action Selection: If person does not have any family member, then he moves to the nearest evacuation location, otherwise the action taken is as shown in Figure 4, where "AllKnown" means the health status of all household members is known.

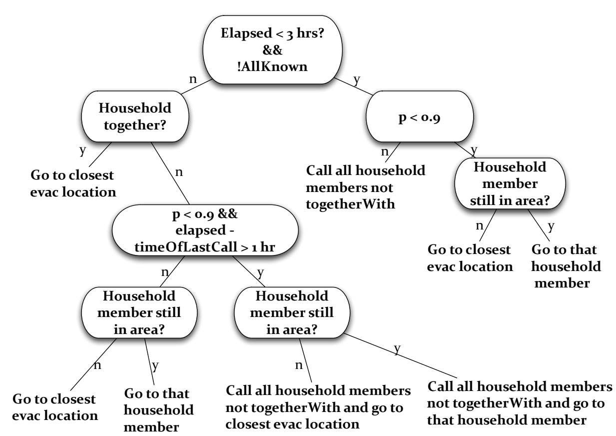

Figure 4: Algorithm for household reconstitution action selection.

Termination Condition: The HRO option is terminated if all family members are at the same location (i.e., they have successfully reconstituted their household). It is terminated with probability 0.5 if somebody in the group is in poor health (*healthstate* < 5, indicating moderate injury or worse), or if it is known that all family members are safe or if somebody in the group has received an emergency broadcast.

5.2.2 *Evacuation (Evac)*

This models individual behavior to evacuate the affected region and move to a safe area. Evacuation destinations are chosen to be the points on major highways just outside the DSA.

Action Selection: An agent attempts to move to the nearest evacuation location. In addition, with probability 0.5, an agent also tries to call all his family members who are not together with him every hour.

Termination Condition: If a person is in poor health ($healthstate < 5$) or is unable to move then this option is terminated.

5.2.3 *Shelter-seeking (Shelter)*

This represents individual behavior of staying inside a building in order to shelter from radiation. A building is designated a shelter location if it has less than 10% damage, though this may not correspond to an equal reduction in exposure as compared to being outdoors, since the reduction in exposure depends on the construction materials of the building and other factors. Detailed data about building conditions were available to us and were incorporated into the model.

Action Selection: If an agent is in a location that provides shelter then it stays there otherwise it tries to move to the nearest shelter location.

Termination Condition: An agent terminates this behavior if its health state falls below 5, or with probability $(1-r)$, where r is percentage radiation attenuated by being inside a building at this location. Some agents are not patient enough to remain in shelter for a long period of time and hence with probability 0.1, they randomly terminate this option.

5.2.4 *Healthcare-seeking (Health)*

This models behavior of a sick or injured agent to seek health care facility.

Action Selection: If an agent is unable to move, it calls 911. If the last call to 911 was successful, then the agent is "teleported" to the nearest health care location with probability 0.2 to mimic being rescued by an ambulance (this probability is low as initially most casualties come to hospital by private vehicle [2]). Otherwise the agent moves toward the nearest healthcare location.

Termination Condition: If an agent is unable to move and all calls attempted fail, then the option is terminated. Otherwise, if the agent is rescued by somebody in "aid & assist" option, or the agent reaches a hospital then this option is terminated.

5.2.5 *Panic*

Though the existence of panic is disputed in the literature [17], it may be expected that in a disaster of this magnitude some people would panic and may behave in a counter-productive or irrational manner.

Action Selection: The action performed by an individual in panic mode is as shown in fig. 5 where "callFlag" means that with probability 0.7, the agent has chosen to call 911, "goOutside" means that with probability 0.5, the agent has chosen to run outside the building (this probability is 0 if the agent is too injured or sick to move), "goHospital" means that with probability 0.3, the agent attempts to move towards a hospital (regardless of its health state).

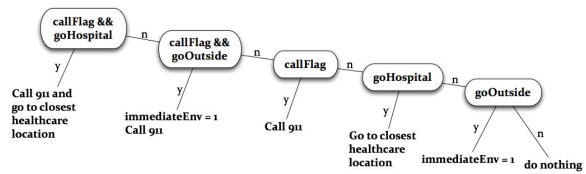

Figure 5: Decision tree for panic action selection.

Termination Condition: People are assumed to be more likely to panic initially than later, hence if the time elapsed since the event is more than 3 hours then with probability 0.5 the agent quits the panic option. To avoid a sharp transition where everyone stops panicking at once, we use a sigmoid function to smooth the probability of quitting the panic option around the 3 hour mark. Alternatively, if an agent has received an emergency broadcast or has made a successful call to 911 then it is less likely to panic, and hence quits the panic option with probability 0.75 and 0.5, respectively.

5.2.6 *Aid & Assist (A&A)*

This models the (survivor) behavior of assisting children (age<5 years), sick people or individuals who are unable to move due to injury.

Action Selection: The algorithm is shown in Figure 6.

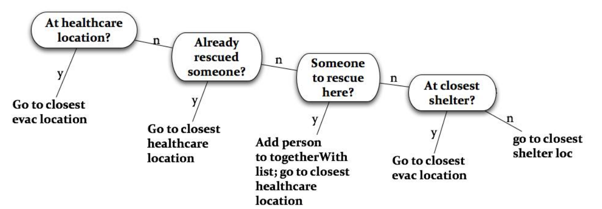

Figure 6: Decision tree for aid and assist action selection.

Termination Condition: An agent quits this behavior if it is sick, unable to move, or somebody in its family is not safe. An agent also quits this option if it is unable to find another agent to rescue at the current location.

5.3 Health Modeling

The health of a person will drive their behaviors, affect their mobility, and influence the time needed for healthcare. For these purposes a simple model that represents health on a continuum for injury triage (based on the SALT triage [1]) is used as the main health state. This continuum consists of states from 0 (death) to 7 (full health), with states 4 and lower corresponding to moderate injury or worse. Secondary effects of health (mobility, health care requirements, etc.) will be calculated based on this state. Additionally, processing individuals for treatment and calculating their response to treatment is also based on these states.

Initial injuries and their severity are calculated based on the physical properties of the blast itself, which has been extensively calculated. If the agent is outdoors, injuries can occur from the physical effects of the blast, which have been modeled to account for the effects of shielding from buildings. Similarly, individuals inside buildings are affected by physical effects of the blast, but can also be injured due to building collapse. Radiation from the blast is also attenuated by buildings and is absorbed by individuals, though the effects of this prompt radiation on the agent's health can be delayed over time.

Following the immediate effects of the blast, an agent's health can deteriorate as a function of time, cumulative radiation exposure, or from injuries suffered while moving over the damaged landscape. Health can be improved for agents who receive healthcare, or as a function of time (mainly for minor injuries). The delayed effects for the prompt radiation exposure from the blast are accounted for as they begin to manifest (for instance absorbing 2.5 Grays of radiation may induce a deterioration of health within 1-4 hours that may impede mobility). Similarly, physical injuries that go untreated for prolonged periods of time can cause delayed deteriorations in health. The likelihood of suffering a health changing injury are based on the physical attributes of the locations a person moves over, those with greater amounts of debris from collapsed buildings etc. are more likely to produce injuries.

The mobility of the individual depends on the severity of the injury, and the likelihood that the injury would prevent an agent from being able to walk (e.g., a broken leg impedes mobility, whereas a broken arm does not). If an agent is severely injured (health state 3) they are very unlikely to be mobile (90% are immobile).

For agents seeking healthcare, they initially seek it at DC area hospitals. As mobile healthcare units brought in by the federal government arrive agents that can see these locations and need health care will seek it there instead of the hospitals. The number of agents that can be treated and the degree of injuries that can be treated depend on the number of health care workers at each facility and the type of facility. Extremely mobile emergency response vehicles are the first to arrive within hours but can deliver very minimal care and larger mobile hospitals take days to arrive but can deliver a wider spectrum of care. Demand for healthcare quickly outpaces the rate, which depends the severity of the injury and number of healthworkers available, at which it can be provided and a queue develops. Agents begin to leave the queue if their injuries are not severe (health state 3) and the line is longer than 10000 agents.

6. EVALUATION

We present initial results from a worst-case analysis. The standard recommendation is for people to seek shelter in the event of a nuclear explosion, in an attempt to minimize radiation exposure [22]. However, it is extremely unlikely that people will be able to determine, on their own, that the event is a nuclear explosion, and so we assume that, in the absence of information, the probability that people will seek shelter is very low [9]. It has been suggested that, through proper education and information dissemination, people can be persuaded to shelter-in-place [11]. For the worst-case analysis, we assume that even on being advised, via emergency broadcasts, to take shelter, the probability of shelter-seeking remains very low. However, we assume that emergency broadcasts have the secondary effect of reducing panic, so that people are more likely to switch to other behavioral options like household reconstitution and evacuation.

This is illustrated in the option-selection decision tree (fig. 3), which shows that $Prob(shelter|EBR) = p$, a parameter (EBR stands for Emergency Broadcast Received). In our simulations, we set $p = 0.1$ so that $Prob(shelter|EBR) = Prob(shelter| \sim EBR)$.

We ran a two cell experiment where we varied the amount of communication restoration. Each cell consists of five independent simulation runs, which we refer to as replicates. Each time step is referred to as an iteration. Each replicate is run for 100 iterations. The first six iterations correspond to 10 minutes of simulated time each, and the remaining correspond to 30 minutes of simulated time each, so that 100 iterations correspond to a total of 2 days simulated time. Since radiation levels vary sharply in the first hour, smaller time intervals were simulated for the first hour.

In the first cell of the experiment, we assume that regions that lose mobile phone coverage do not regain it for the duration of the simulation. In the second cell, we assume that mobile phone coverage cannot be restored within 0.6 miles of ground zero, but is restored to 50% capacity within 3 hours in the 0.6 to 1 mile ring. Outside the 1 mile ring, coverage remains at full capacity in both cells of the experiment.

Being able to communicate affects agent behavior in multiple ways. Receiving emergency broadcasts and making successful 911 calls both help to reduce panic. This makes more agents switch to other behavioral options early in cell 2. If agents are able to determine that their family members are safe, then they are also more likely to switch to the Aid & Assist option. Figure 7 shows the difference in the average number of agents executing each behavioral option in each iteration. The counts are averaged over five independent runs in each cell.

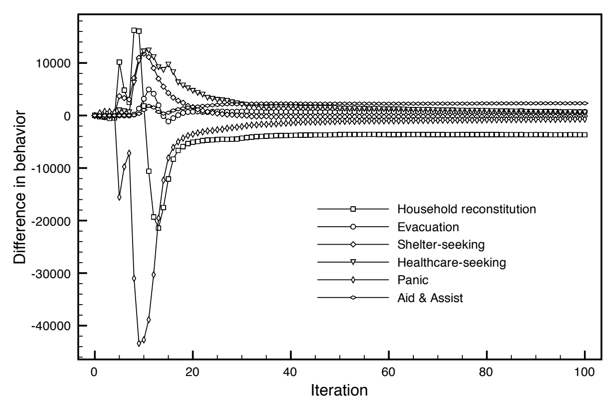

Figure 7: This figure shows the difference in the number of agents executing each behavior in each iteration. The counts are averaged over five independent runs, and the difference is cell 2 (partial communication restoration) - cell 1 (no restoration).

These differences in behavior manifest a difference in the number of injured people in the two cells. Figure 8 shows the differences in the number of agents with low healthstate (< 5, i.e., moderate injury or worse) over time (iterations). We see that initially, agents in cell 2 are worse off, but over time the difference turns in favor of cell 2. To explain this difference as resulting from a combination of behaviors, we do a linear regression to fit the values in fig. 8 using the behavior differences in fig. 7. The result is shown as a solid line in fig. 8.

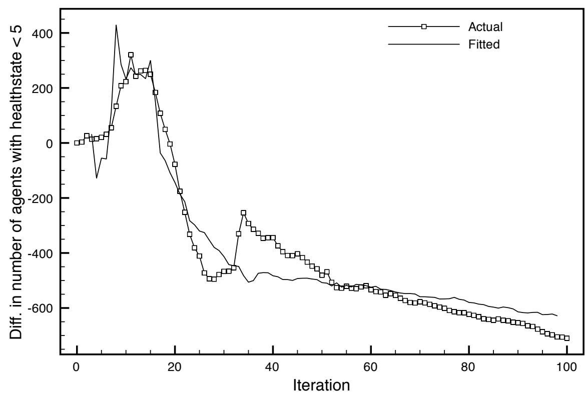

Figure 8: This figure shows the difference in the number of agents with $healthstate < 5$ (moderate injury or worse) in each iteration. The counts are averaged over five independent runs in each cell. We see that the difference is increasingly negative over time, indicating that agents in cell 2 (partial communication restoration) are in better health. The fitted curve (solid red line) is calculated as a linear regression of the behavior differences from fig. 7. The coefficients are shown in table 2.

The coefficients and significance values are shown in table 2. We see that all the behaviors are significant, though they have different effects. The largest contribution is due to the aid & assist behavior, and its negative coefficient indicates that this behavior contributes to reducing the number of unhealthy agents in cell 2 (hence making the difference more negative in fig. 8). The sheltering behavior has a similar effect, though its magnitude is slightly less than half that of aid & assist.

On the other hand, the largest positive coefficients are for healthcare-seeking and evacuation behaviors, which implies that both these behaviors actually have a deleterious effect on health.

Household reconstitution and panic have statistically significant effects also, but their magnitudes are much smaller than the effects due to the other behaviors.

Table 2: Coefficients for the linear regression shown in fig. 8.

	Estimate	Std. error	t value	Pr($>\|t\|$)
(Intercept)	-232.3	68.03	-3.415	0.000962
HRO	0.02575	0.00586	4.394	3.06e-5
Evac	0.1004	0.002658	3.776	0.000287
Shelter	-0.09226	0.02098	-4.397	3.03e-5
Health	0.1267	0.01546	8.195	1.73e-12
Panic	-0.01032	0.003295	-3.132	0.002351
A&A	-0.2014	0.03474	-5.797	1.01e-7

We can see, from fig. 7, that the numbers of agents attempting to evacuate and to seek healthcare early in the simulation are higher in cell 2. This early rush to leave the affected area actually causes increased exposure as well as increased injury for agents in cell 2, which causes the early peak in fig. 8. This explanation is supported by data in fig. 9, which shows the difference in exposure and injury levels between cells 2 and 1, both of which peak early also. The difference in exposure seems small, only up to 0.2 cGy, but that is due to the fact that it has been averaged over the entire population. The difference is much higher for agents who start out close to ground zero, though this is not shown for lack of space.

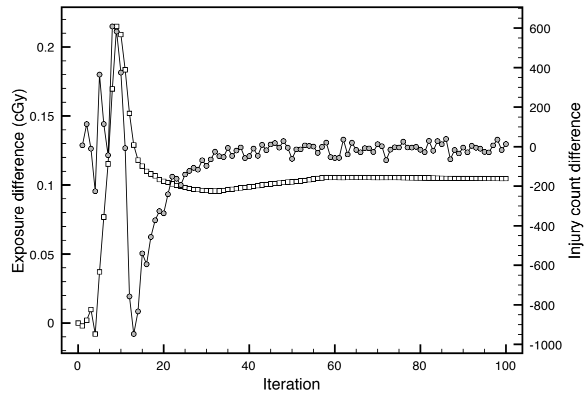

Figure 9: This figure shows the difference in exposure levels (squares) and the difference in injury counts (filled circles) between cells 2 and 1. Both the curves show an early peak before settling into a steady state.

However, after a little while, the benefits of aiding & assisting kick in and help to reduce the number of unhealthy agents in cell 2. Note that this is an indirect benefit of restoring communication in two ways. Some agents are able to determine that their family members are safe by being able to call/text them, and thus switch to the aid & assist behavior. Some other agents are able to find their family members sooner even if they aren't able to contact them (because they stop panicking early upon receiving emergency broadcasts or making successful 911 calls), and subsequently switch to the aid & assist behavior if all the family members are healthy.

7. CONCLUSION

This paper presents aspects of agent design and agent interaction design of a large multiagent simulation study of a hypothetical improvised nuclear detonation in Washington DC.

We have constructed a detailed model of a behaving human population using a synthetic population model for Washington DC, augmented with a model of human behavior based on a decentralized SMDP model. The set of behaviors represented are chosen from the literature and surveys about anticipated public response in such a scenario.

High levels of prompt radiation and fallout will make it impossible to carry out immediate relief operations. Therefore, it is recommended that people shelter-in-place for up to 12 hours after the explosion [22]. We have done a worst-case analysis where, even upon being advised to do so, the probability that people actually shelter-in-place remains low. Our analysis shows that a relatively passive intervention, partially restoring communication capacity, can have a significant positive effect on people's health.

This beneficial effect emerges from the complex interactions between human behavior and infrastructure. Increased communication has the effect of reducing panic and increasing movement as people try to find their families and to evacuate. This has the adverse effect of causing increased exposure and injury. However, as people are able to establish the health and safety of their family members, they also turn to aiding and assisting others in greater numbers, which in turn has a positive health effect.

Much work can be done to extend this study. Previous studies were restricted to asking questions about infrastructural damage and estimating casualties and mortalities based on static landscan data. By having a model of a behaving population, we are now in a position to ask new questions, such as how can we reduce various measures of distress (reuniting people with their families, evacuating people and returning them to their homes, etc.), how can we optimally place relief resources like mobile healthcare units, and how can we optimally deploy resources like firefighting units and police units. The important change of perspective in our work is to show that human instincts and responses actually present opportunities to shape behavior by relatively passive interventions, and thereby help people help themselves.

Acknowledgements

We thank our external collaborators and members of the Network Dynamics and Simulation Science Lab (NDSSL) for their suggestions and comments. This work is supported in part by DTRA CNIMS Contract HDTRA1-11-D-0016-0001, DTRA Grant HDTRA1-11-1-0016, NSF NetSE Grant CNS-1011769, NSF SDCI Grant OCI-1032677, and NIH MIDAS Grant 2U01GM070694-09.

8. REFERENCES

[1] SALT mass casualty triage. *Disaster Medicine and Public Health Preparedness*, 2(4):245–246, Dec 2008.

[2] E. Auf der Heide. The importance of evidence-based disaster planning. *Annals of Emergency Medicine*, 47(1):34–49, Jan 2006.

[3] C. Barrett, S. Eubank, A. Marathe, M. Marathe, Z. Pan, and S. Swarup. Information integration to support policy informatics. *The Innovation Journal*, 16(1):article 2, 2011.

[4] R. J. Beckman, K. A. Baggerly, and M. D. McKay. Creating synthetic baseline populations. *Transportation Research Part A: Policy and Practice*, 30(6):415–429, November 1996.

[5] B. R. Buddemeier, J. E. Valentine, K. K. Millage, and L. D. Brandt. National Capital Region: Key response planning factors for the aftermath of nuclear terrorism. Technical Report LLNL-TR-512111, Lawrence Livermore National Lab, November 2011.

[6] M. J. Dombroski and P. S. Fischbeck. An integrated physical dispersion and behavioral response model for risk assessment of radiological dispersion device (RDD) events. *Risk Analysis*, 26(2):501–514, 2006.

[7] T. E. Drabek and K. S. Boggs. Families in disaster: reactions and relatives. *Journal of Marriage and the Family*, 30:443–451, 1968.

[8] C. V. Goldman and S. Zilberstein. Communication-based decomposition mechanisms for decentralized MDPs. *J. Artif. Int. Res.*, 32(1):169–202, May 2008.

[9] T. M. Guterbock, J. H. Lambert, R. A. Bebel, and M. W. Parker. NCR behavioral survey 2011: Work, school or home? Issues in sheltering in place during an emergency. Technical report, Center for Survey Research, University of Virginia, Aug. 2011.

[10] R. D. Lasker. *Redefining Readiness: Terrorism Planning Through the Eyes of the Public.* Center for the Advancement of Collaborative Strategies in Health, New York Academy of Medicine, 2004.

[11] R. D. Lasker, N. D. Hunter, and S. E. Francis. *With the Public's Knowledge, We Can Make Sheltering in Place Possible.* New York Academy of Medicine, New York, NY, 2007.

[12] Y. Lin, I. Fedchenia, B. LaBarre, and R. Tomastik. Agent-based simulation of evacuation: An office building case study. In W. W. F. Klingsch, C. Rogsch, A. Schadschneider, and M. Schreckenberg, editors, *Pedestrian and Evacuation Dynamics 2008*, pages 347–357. Springer Berlin Heidelberg, 2010. 10.1007/978-3-642-04504-2_30.

[13] S. Liu, P. Murray-Tuite, and L. Schweitzer. Analysis of child pick-up during daily routines and for daytime no-notice evacuations. *Transportation Research Part A: Policy and Practice*, 46(1):48–67, Jan. 2012.

[14] X. Pan. *Computational Modeling of Human and Social Behaviors for Emergency Egress Analysis.* PhD thesis, Dept. of Civil and Environmental Engineering, Stanford University, 2006.

[15] N. Parikh, S. Shirole, and S. Swarup. Modeling the effects of transient populations on epidemics. In *The AAAI Fall Symposium on Social Networks and Social Contagion*, Arlington, VA, Nov 2012.

[16] N. Pelechano, K. O'brien, B. Silverman, and N. Badler. Crowd Simulation Incorporating Agent Psychological Models, Roles and Communication. In *1st Int'l Workshop on Crowd Simulation*, pages 21–30, 2005.

[17] R. W. Perry and M. K. Lindell. Understanding citizen response to disasters with implications for terrorism. *Journal of Contingencies and Crisis Management*, 11(2):49–60, 2003.

[18] M. F. Sherman, M. Peyrot, L. A. Magda, and R. R. M. Gershon. Modeling pre-evacuation delay by evacuees in World Trade Center towers 1 and 2 on September 11, 2001: A revisit using regression analysis. *Fire Safety Journal*, July 2011.

[19] B. R. Steunebrink, M. Dastani, and J.-J. C. Meyer. Emotions to control agent deliberation. In *Proc. AAMAS*, pages 973–980, Richland, SC, 2010.

[20] R. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1-2):181–211, 1999.

[21] J. Tsai, N. Fridman, E. Bowring, M. Brown, S. Epstein, G. Kaminka, S. C. Marsella, A. Ogden, I. Rika, A. Sheel, M. Taylor, X. Wang, A. Zilka, and M. Tambe. ESCAPES - Evacuation simulation with children, authorities, parents, emotions, and social comparison. In *Proc. AAMAS*, Taipei, Taiwan, May 2011.

[22] L. M. Wein, Y. Choi, and S. Denuit. Analyzing evacuation versus shelter-in-place strategies after a terrorist nuclear detonation. *Risk Analysis*, 30(6):1315–1327, 2010.