

PLANNING AND RESPONSE IN THE AFTERMATH OF A LARGE CRISIS: AN AGENT-BASED INFORMATICS FRAMEWORK

Christopher Barrett, Keith Bisset, Shridhar Chandan, Jiangzhuo Chen, Youngyun Chungbaek,
Stephen Eubank, Yaman Evrenosoğlu, Bryan Lewis, Kristian Lum, Achla Marathe,
Madhav Marathe, Henning Mortveit, Nidhi Parikh, Arun Phadke, Jeffrey Reed, Caitlin Rivers,
Sudip Saha, Paula Stretz, Samarth Swarup, James Thorp, Anil Vullikanti, Dawen Xie

Departments of Computer Science, Agriculture and Applied Economics,
Electrical and Computer Engineering,
Network Dynamics and Simulation Science Laboratory
Virginia Bioinformatics Institute
Virginia Tech
Blacksburg, VA 24061, USA

Contact Authors: Madhav Marathe (mmarathe@vbi.vt.edu) and Samarth Swarup (swarup@vbi.vt.edu)

ABSTRACT

We present a synthetic information and modeling environment that can allow policy makers to study various counter-factual experiments in the event of a large human-initiated crisis. The specific scenario we consider is a ground detonation caused by an improvised nuclear device in a large urban region.

In contrast to earlier work in this area that focuses largely on the prompt effects on human health and injury, we focus on co-evolution of individual and collective behavior and its interaction with the differentially damaged infrastructure. This allows us to study short term secondary and tertiary effects. The present environment is suitable for studying the dynamical outcomes over a two week period after the initial blast.

A novel computing and data processing architecture is described; the architecture allows us to represent multiple co-evolving infrastructures and social networks at a highly resolved temporal, spatial, and individual scale. The representation allows us to study the emergent behavior of individuals as well as specific strategies to reduce casualties and injuries that exploit the spatial and temporal nature of the secondary and tertiary effects.

A number of important conclusions are obtained using the modeling environment. For example, the studies decisively show that deploying ad hoc communication networks to reach individuals in the affected area is likely to have a significant impact on the overall casualties and injuries.

1 INTRODUCTION

Large-scale natural and human-initiated crises that affect significant population centers are important from the standpoint of national security, resilient design of cities, and the overall security and well being of society. Hurricanes present the most commonly occurring natural disasters in the US. The effects of hurricane Sandy that hit New York City provide a clear demonstration of the short as well as long term impacts of such disasters.

The focus of this paper is on human-initiated crises, although many of our broad conclusions are likely to be applicable in more general settings. Unlike natural disasters, our understanding of the effects of large-scale human initiated crises is limited; we are fortunate not to have experienced too many such events as global citizens. Nevertheless, planning for such scenarios is vitally important in order to be prepared should such an event actually occur. Computer simulations and an associated modeling and informatics

environment provide a natural analytical tool that can aid in developing and assessing various policies and response plans in the event of such crises. Here we focus on a hypothetical scenario published by the Department of Homeland Security, which is essentially similar to a scenario that has been studied for the past 60 years. In this scenario, an improvised nuclear device (IND) is detonated in a large urban region.

1.1 Contributions

In this work we focus on the secondary and tertiary effects of the IND detonation. We are interested in the role of adaptive human behaviors and their co-evolution with the damaged infrastructure. The paper describes a comprehensive and integrated modeling and decision support environment that facilitates the study of various counter-factual scenarios related to the aftermath of such an event. We leverage the earlier work on calculating various prompt effects related to fallout (Buddemeier, Valentine, Millage, and Brandt 2011). Since our focus is on individuals and co-evolving behaviors, we have developed a highly resolved synthetic representation of the National Capital Region that includes Washington DC, as described in Section 4. Briefly, it represents the dynamic synthetic populations and urban built infrastructure including buildings, power, transportation, health, and communication infrastructures, and a description of individual and family behaviors.

Our results show that detailed behavioral representation and its co-evolution with the physical infrastructure has subtle but important effects and implications. They provide details on spatial and temporal distribution of populations, their behaviors and affective states. The representation allows us to represent and track important sub-populations, e.g., elected federal officials, critical workers, etc., allowing one to understand contextual information and yet leading to the formulation of general guidelines.

A novel aspect of our work is the synthetic information and modeling environment that we have developed for representing and studying the problem. The architecture overcomes traditional problems associated with HLA (High level architecture) used in the traditional DoD literature. A key concept developed and refined here is that of *un-encapsulated agent representation*: the representation of individual behavior as a composition of high level intents and goals with refinements that are specific to infrastructures. Publicly available datasets (open source as well as commercial) combined with expert knowledge (abstracted as procedural information) were used to synthesize infrastructure networks, synthetic populations, and individual and component level behavior. Another novel aspect of our work is the use of a relational database to drive high performance computing oriented dynamic models. Our work is motivated by the earlier work by Gray and his colleagues (Heber and Gray 2007a, Heber and Gray 2007b), but extends it to drive dynamic HPC-oriented agent models.

The synthetic information and modeling environment allows us to estimate every individual's mobility patterns after the event. This information allows us to align response policy with survivors' behaviors and needs to improve compliance. It allows us to study strategies for locating triage centers, field hospitals and other logistical and medical support installations. The detailed representation of the population allows us to better estimate fallout exposure and quantify the relevant statistics for specific demographic groups. It allows us to study prioritizing and staging response and recovery efforts, e.g., whether it is more important to restore communications or to send emergency responders within the fallout zone. Finally, spatiotemporal demand generated by survivors for essential resources such as water, electricity and communication can be better estimated, leading to a better understanding of the inherent electrical power, communications, transportation and other emergency support functions. Important insights based on the counter-factual experiments include:

1. The role of the wireless communication network. A small improvement in communication networks has a disproportionately large and positive impact on the overall behavior, leading to fewer deaths, better health outcomes and reduction in panic. Recent progress in communication technology, some of which was motivated by earlier disasters, includes: pervasive availability of mobile phones, the ability of the cell towers to withstand an electro-magnetic pulse (EMP) at moderate distance and the ability to deploy mobile cellular infrastructure such as COWS (cells on wheels) or "network in a suitcase" plays a critical role in

reaching people via SMS, voice and digital messages and emergency broadcasts. Additionally, mobile phones are now carried by a large fraction of the population. This allows individuals to be reached even when they are not at a fixed location.

2. The role of individual behavior. Despite being a largely physical event, it is important to represent human behaviors and their adaptation. Large-scale spatiotemporal patterning of behavior that emerges as a result of the interactions between individual citizens and infrastructures has important consequences on the efficacy of public policies. For example, family reconstitution is an important task that many citizens are likely to undertake in the aftermath of a large scale disaster (Lasker 2004). This is likely to be true even if the government advises citizens to seek shelter to avoid exposure to harmful radiation. Transport and communication networks will be burdened by these actions. The emergent collective behavior depends on the individual behavior and adaptation.

3. Which behaviors to represent. The relevant set of behaviors changes over time. Panic is likely to occur early, while search and rescue efforts will take several hours to get under way. Similarly, initially people will be concerned about finding their families, evacuating the area, etc., but later there will be the problem of relocating the displaced population. Statistical analysis of simulation results can tell us when the represented behaviors are relevant, and at what point the set of behaviors represented needs to be changed or updated.

4. Disruption to the power network and its long term impact. The power network suffers a huge loss and large portions of the network will likely be inoperational for at least a year or two. This will have important implications for how the city and its surroundings will be reconstituted. This suggests that the country might benefit from developing a power network infrastructure stockpile comprising of network elements, transformers etc.

2 SCENARIO

The basic scenario is described as the National Planning Scenario 1 in a document published by the Dept. of Homeland Security. It has been discussed in a number of earlier articles and thus we will describe the scenario briefly. Additional information can be found in (Federal Emergency Management Agency 2010, Wein et al. 2010). The scenario we model is based on previous models of the physical effects of the detonation (Federal Emergency Management Agency 2010, Buddemeier et al. 2011, Wein et al. 2010). This sort of scenario, focusing on the physical effects, has been studied for 60 years now. However, none of the previous work takes human behavioral response into account.

A hypothetical low yield (10kT) nuclear device is detonated at ground level in Washington DC (an urban region). Specific location and time of this *hypothetical* event is: -77.036574 longitude, 38.902604 latitude, Washington DC Time: 11:15 EST Date: May 15, 2006, which is at the intersection of 16th and K streets. In the present literature this is called an improvised nuclear device (IND). Our focus is on a region we call the detailed study area (DSA), shown in Figure 1a, which is the area defined by the .01 Gy fallout contour at 60 minutes joined with the thermal radiation contour at 2.1 cal/cm² bounded by the boundary of the counties neighboring the District of Columbia.

The physical damage to the infrastructure, fallout and radiation levels, blast effects, etc. were provided to us and are shown in Figure 1b. Critical infrastructure in the Washington DC region is damaged by the IND; roadways are filled with rubble; cell towers within 0.6 miles of ground zero (GZ) are destroyed, a large area around GZ suffers a long term blackout. Most buildings within 1000 meters of the detonation are severely damaged. The EMP destroys communication networks within ~3 miles of GZ. Intense heat causes numerous fires. The immediate effect on human health is as follows: Deaths: 279K; Injured 93K. These numbers are in relatively good correspondence with previous analyses, although these are bottom-up, individual behavior-based, individualized injury calculations.



Figure 1: The detailed study area and the damage caused by the blast, along with the fallout cloud.

3 PREVIOUS WORK

A number of researchers have studied this scenario in detail. Notable work includes the work of Buddemeier, Valentine, Millage, and Brandt (2011), Wein, Choi, and Denuit (2010), and Federal Emergency Management Agency (2010). Wein et al. study the trade-off between evacuation and shelter-in-place and show that shelter-in-place is an extremely important policy from the standpoint of reducing the secondary and tertiary health effects. Buddemeier et al. also make a similar case for shelter-in-place. Their work, which has been done for over a decade now, is comprehensive in many aspects and has been published as a series of progressively more detailed reports. An important and distinguishing feature of the work done by Buddemeier et al. is a detailed representation of the buildings and calculations of various prompt effects. The resulting calculations provide good estimates of the overall health and injury levels caused by the IND explosion.

Our synthetic information-based modeling and simulation infrastructure goes beyond these studies in its ability to represent a behaving population. A number of researchers have developed general purpose platforms for agent-based simulations; see (Collier and North 2011, Kin and Chan 2010, Hybinette et al. 2006, Riley and Riley 2003) and the references therein for further details. The high performance computing oriented modeling environment we have developed is not as general as systems like Repast (Macal and North 2010, North et al. 2013) or Sassy, but it scales adequately for the restricted class of problems we study. It also allows detailed representations of various infrastructures. For example, Repast enables agents to interact in two distinct kinds of environments – one in which agents are connected in a network, the other wherein agents interact over a spatially explicit grid. Agents in our modeling framework interact via a network that is often spatially explicit. Furthermore, agents also interact with fixed infrastructure elements. Perhaps more interestingly, agent interaction is mediated via a number of different networks; each network representing a specific infrastructure system. The overall representation of the agent thus is obtained by composing the infrastructure specific behaviors with a set of generic behaviors/affects that are in a sense infrastructure independent (e.g., panic). Section 4.1 provides details of the computational models and the high-performance computing-based architecture we have developed.

4 SYNTHETIC INFORMATION, INFRASTRUCTURES, AND NETWORKS

A key contribution of the work described here is to illustrate the relevance of detailed social and individual behavior, even in such a physically dominated event. This requires that we develop detailed yet synthetic representations of the urban region under consideration.

A synthetic information system in our context comprises of synthetic individuals, infrastructure elements and networks as well as behavioral representations of the agents and the infrastructure. We call it synthetic information systems since: (i) such information is synthesized from diverse sources of information that is often available at various levels of aggregation, (ii) individuals comprising the populations are not real but

Table 1: Datasets for the synthetic population and infrastructure used in the simulation.

(a)		(b)	
Synth. pop.	Data sources	Infrastructure	Data sources
Base US population	Amer. Community Survey, TIGER/Line shapefiles, Nat. Center for Education Stat., Nat. Household Travel Survey, Navteq, Dun & Bradstreet	Transportation	Navteq, WMATA
Transient pop. (additional)	Destination DC, Smithsonian visit counts	Communication	TowerMaps, Sprint & AT&T API, CDC TUS
Dorm students (additional)	CityTownInfo, DC public access - online Data Catalog	Power	Pepco, DC govt, Google Earth, PSSE (Siemens)
		Healthcare	Dun & Bradstreet, National Registry of Hospitals

based on aggregate data and thus are statistically indistinguishable from the real population. Both these aspects are important. Furthermore, it is important that individual and collective behaviors are an integral part of the synthetic information. This is procedural data and is an important component when developing dynamic models.

Synthetic populations and social networks are created based on an extension of methodology we have developed over the last 15 years. Unlike simple random graph techniques, these methods use real world data sources and combine them with behavioral and social theories to synthesize networks and populations of interest. A key aspect of such synthetic networks is the fact that they possess a rich structure that goes beyond simply being scale-free networks. This rich structure in turn affects both the computational efficiency of the algorithms to process the networks as well as the dynamical outcomes over such networks. See (Gonzalez et al. 2008, Barrett et al. 2009, Beckman et al. 2013) for additional discussion.

We develop a synthetic population for the Washington DC urban region that models every individual in the population. The synthetic population is a set of synthetic people and households, geographically located, each associated with demographic variables drawn from any of the demographics available in the census. Household structure and demographics are derived from U.S. Census data. Each synthetic individual is assigned a 24-hour activity sequence including geo-locations for each activity. A social contact network is constructed based on physical co-location of the interacting persons. Our work builds on our earlier work in synthesis and analysis of large relational networks. Initial work was done under the TRANSIMS and NISAC projects and more recently new methods have been developed under the CNIMS project (Barrett et al. 2009, Beckman et al. 1996, Barrett et al. 2001). The process preserves the confidentiality of the original data sets, yet produces realistic attributes and demographics for the synthetic individuals. A substantial effort has been spent on calibration and validation of our synthetic populations and networks; see (Barrett et al. 2001, Chowell et al. 2003, Eubank et al. 2004, Barrett et al. 2007, Beckman et al. 1996) for details. First, the design of the system is based on a formal theory of simulation called Sequential Dynamical Systems (Eubank et al. 2006, Barrett et al. 2003, Barrett et al. 2007). Various microscopic and macroscopic quantities produced by TRANSIMS have been validated in the city of Portland, including (i) traffic invariants such as flow density patterns and jam wave propagation; (ii) macroscopic quantities, such as activities and population densities in the entire city, number of people occupying various locations in a time varying fashion, time varying traffic density split by trip purpose and various modal choices over highways and other major roads, turn counts, number of trips going between zones in a city, etc. Results on population mobility and social network construction were presented and reviewed annually at (Barrett et al. 2001).

The synthetic population for the Washington DC metro area (which includes surrounding counties) is used to initialize the DSA. The daily activities of the population cause some people who reside in the DSA

to be out of the area when the detonation happens, and cause some people who reside outside the area to be within the DSA at the time. This also means that most people have family members distributed both inside and outside the DSA.

A summary of the infrastructures and populations represented in our modeling environment is given in Table 1. In addition to the population and their locations at the time of the detonation, the initialization conditions also include the effect of the blast, including damage to the transportation, power, and communication infrastructures. The damage to the power infrastructure is not recoverable in the duration of the simulation. The transportation and communication infrastructures have a role to play in the outcomes because they interact with human behavior: the movement of the population can cause congestion which affects movement times and decisions about where to go or what to do. It also stresses the communication infrastructure, especially because many cell towers are inoperative around ground zero. We also model an intervention where Cells-on-Wheels (CoWs) can be brought in to partially restore communication close to ground zero. For these reasons, transportation and communication are modeled dynamically over the course of the simulation, whereas electrical power is not.

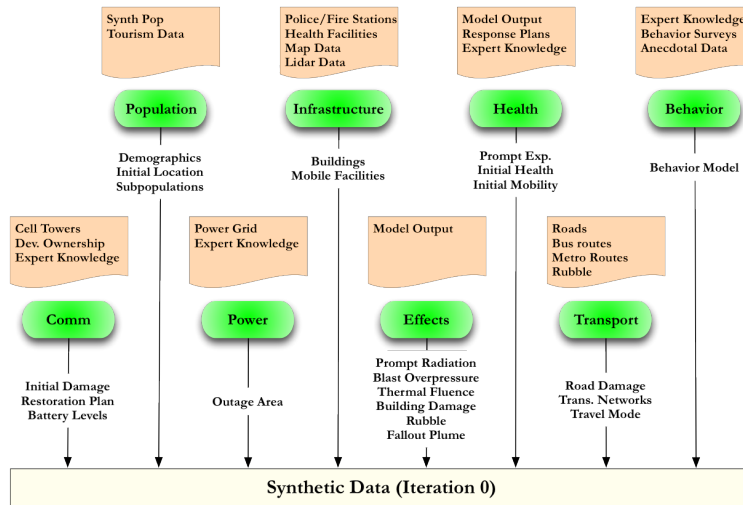


Figure 2: Constructing the initial conditions for the simulation.

Several extensions to the base populations and networks were required for carrying out the present computational experiments. These include: (i) representing transient populations, given that they are a significant fraction in DC; (ii) representing multitude of behaviors and affects as they pertain to the aftermath of an IND explosion (discussed further in our papers (Lewis et al. 2013, Adiga et al. 2013, Chandan et al. 2013, Barrett et al. 2012)); and (iii) developing methods to synthesize intermodal transportation networks (including vehicular and pedestrian transport, metro, buses), electrical power networks and communication networks. Much of the data for the intermodal transport network is available, but substantial effort was spent in integrating these data sets to produce a computer readable and consistent representation. For example, the bus network has to be geographically consistent with the underlying road network. Same holds for the metro. The bus stops and train stops have to be geographically located and the location has to be consistent with respect to the overall spatially explicit system. Communication and electrical power networks pose a more vexing problem. For communication networks and increasingly for electrical networks, very little open source data is available. Companies have cited various reasons including competitive advantage, security, etc for not making such data sets available. We thus had to synthesize the best extant representations of these networks. We were able to do this quite well for electrical power networks. For communication networks, we were able to represent certain elements, including cell towers. Connectivity between these individual elements was established using established engineering principles and expert knowledge. See

(Lewis et al. 2013, Adiga et al. 2013, Chandan et al. 2013, Barrett et al. 2012) for additional details on methods for constructing the context-specific synthetic information system for Washington DC.

Figure 2 shows the initial conditions for our models. Besides the population and the infrastructures, we take into account blast and prompt radiation effects, the fallout cloud (which moves east by northeast-ward over time), the locations of hospitals, shelter and evacuation locations, and damage to each of them.

The simulation proceeds in *iterations*, which correspond to ten minutes of simulated time per iteration for the first hour after the detonation, and a half hour per iteration for the next 47 hours. To simulate two days, therefore, we need 100 iterations of the basic dynamic model.

4.1 HPC-based Scalable Computational Architecture

We use a high performance database-driven architecture for calculating the dynamical outcome of co-evolving multi-networks. The basic approach differs from traditional agent-based models. It also differs from the HLA architectures used in the DoD. Each infrastructure is represented using an extension of an agent-based model. The key extension concerns the use of unencapsulated agents — representation of individuals is not stored in a single object as in traditional agent-based models, but rather distributed across infrastructures. This distributed representation of each agent allows us to scale the system efficiently. Modules are loosely coupled and the coupling is achieved via information exchange. In this sense, any module that can provide appropriate guarantees on the input and output can be used within the framework. We use a relational database to coordinate this data exchange. Data for each iteration are stored in two main tables: a *person* table and a *location* table. Modules corresponding to human behavior and the various infrastructures operate on these tables in turn, each transforming the information for the next module.

An iteration begins with the **communication module**, which updates the availability of the cellphone network at each location in the DSA. The **effects module** then updates the radiation levels and the fallout cloud for each location. These two steps involve pre-computed quantities and do not depend upon outcomes in the previous iteration.

Table 2: Summary of high-level behaviors.

Behavior	High-level description
Household reconstitution	Call, move towards household members
Evacuation	Move outside region
Shelter-seeking	Shelter in place or move towards shelter
Healthcare-seeking	Call 911, move towards hospital
Panic	Call 911, run outdoors, move towards hospital
Aid & Assist	Transport hurt individuals to hospital

The **behavior module** is executed at this stage, which updates the current behavior for each agent, depending on various factors like their location, their health state, their knowledge of family members’ health states, and outcomes of attempted phone calls in the previous time step. There are six different behaviors agents can choose amongst: household reconstitution, evacuation, shelter-seeking, healthcare-seeking, panic, and aid & assist. These are summarized in table 2. An agent can only engage in one behavior at a time; in that sense they are mutually exclusive. However, each behavior consists of a choice of movement (destination) and calling (family members and/or 911). In that sense the behaviors are interrelated. The behavior module is designed as a decentralized semi-Markov Decision Process (dec-SMDP) with communication, using the framework of options (Sutton et al. 1999) (which we are referring to as behaviors). Note that even though the module is designed using an MDP framework, we are not doing any optimization because the goal is to capture natural human behavior (and the scale of the simulation is prohibitive too). However, the use of the formalism aids understanding and offers extensibility and expressiveness. The behavior selection algorithm is detailed in (Parikh et al. 2013).

The **transportation module** is responsible for determining a route for each agent based on its current destination (some agents may not be moving; these are ignored at this step), and for moving the agent (partially) along this route. The speed of an agent depends on its healthstate and the conditions of the roads and walking paths the agent is moving over. This step, therefore, alters the location of each agent that is moving. We also implement group movements, e.g., when family members meet up they move together thereafter. More details of the transportation module can be found in (Adiga et al. 2013).

The communication module is then executed again to determine call success/failure for every call that the agents tried to make (as specified by the behavior module). This step also results in the battery drain of cell phones, which is tracked over the course of the simulation. We also include the possibility of *emergency broadcasts* via the cell phone network, which can be done by the government in order to inform people of what has taken place and to advise them to shelter in place. This module is further detailed in (Chandan et al. 2013).

Finally the **health module** evaluates the change in the healthstate of each agent based on their exposure to radiation during the current iteration, the possibility of injury as they move over the damaged landscape, deterioration of health due to earlier exposure and injuries, and possible improvement of health due to received healthcare. Further details of the healthcare model can be found in (Lewis et al. 2013).


The simulation then moves to the next iteration.

5 STUDIES AND RESULTS

We have done multiple studies to evaluate the effects of communication availability on health-related outcomes. We consistently find that a relatively passive intervention, rapidly restoring partial communication availability to regions close to ground zero, can have a significant impact on saving lives and reducing injuries. We detail one of the experiments below.

Cell 1: No restoration $Pr(shelter EBR) = 0.1$	Cell 2: Partial restoration $Pr(shelter EBR) = 0.1$
Cell 3: No restoration $Pr(shelter EBR) = 0.9$	Cell 4: Partial restoration $Pr(shelter EBR) = 0.9$

(a) Experiment design varying communication restoration and a behavioral parameter.



(b) Comparison of the number of people with moderate injuries or worse.

Figure 3: A four-cell experiment and comparison of health outcomes.

In the experiment, we consider two scenarios for restoration of communication. In the first scenario, we assume regions that lose mobile phone coverage do not regain it for the duration of the simulation. In the second scenario, 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 scenarios.

Second, we assume that emergency broadcasts will be sent out over the cellphone network, advising people to shelter in place. However, people may not actually choose to do so; they may be more concerned about finding their families, for example. To evaluate the effects of population response, we create a probability of sheltering, given than an emergency broadcast is received. This is implemented in the behavior module, and we consider two extreme values, $Pr(shelter|EBR) = 0.1$ and $Pr(shelter|EBR) = 0.9$, where EBR stands for Emergency Broadcast Received. In the behavior module, the probability of choosing the shelter-seeking behavior is 0.1 even if an emergency broadcast is not received, so the first condition

makes essentially no difference to the proportion of people sheltering, though it is assumed to have secondary effects like reducing panic.

Previous studies have claimed that sheltering in place is the best policy to follow for several hours in the aftermath of the detonation (Wein et al. 2010). However, intuitively, many people will not know that a nuclear device has been detonated and will not know to shelter in place. It has also been claimed that the annular region between 0.6 miles and 1 mile from ground zero is the region where there is the greatest opportunity to save lives (Buddemeier et al. 2011) if people take shelter as early as possible. People closer to ground zero are expected to have a very low probability of survival and people further out are expected to have a high probability of survival. Our experiment is therefore designed to quantify these effects by restoring communication in this area and sending emergency broadcasts advising people to shelter in place.

Together, the two infrastructure scenarios and the two behavioral probabilities give us a four cell experiment, as shown in Figure 3a. We did five simulation runs, termed replicates, for each cell. Each replicate was run for 100 iterations (time steps). Figure 3b shows a comparison between cells of the number of people with moderate injuries or worse over the course of the simulation. Differences are calculated by first taking the mean of the five outcomes within each cell, and then subtracting the mean for the other cells as indicated. We note a few interesting observations about the outcomes:

- 1) The difference between cells 3 and 1 is close to zero. This means that if communication is not restored close to ground zero, emergency broadcasts will not have any effect on health outcomes.
- 2) Cell 2 does better than cell 1 even though $Pr(shelter|EBR)$ is the same for both. This means that, even though the motivation for restoring communication might be to send emergency broadcasts, communication restoration has a positive impact on health even if emergency broadcasts are not sent.
- 3) If communication is partially restored, and $Pr(shelter|EBR)$ is high, then we get a large benefit, as shown by the comparison between cells 4 and 3 in Figure 3b. Note that to evaluate this difference, we should really compare cells 4 and 1, but since cells 3 and 1 result in almost no difference, we have just shown the comparison between cells 4 and 3.

The first observation above validates the claim from the previous study that the opportunity to save lives is mostly in the region from 0.6 to 1 mile from ground zero. To gain insight into observations 2 and 3 above, we look at the difference in behaviors between the cells over the course of the simulation.

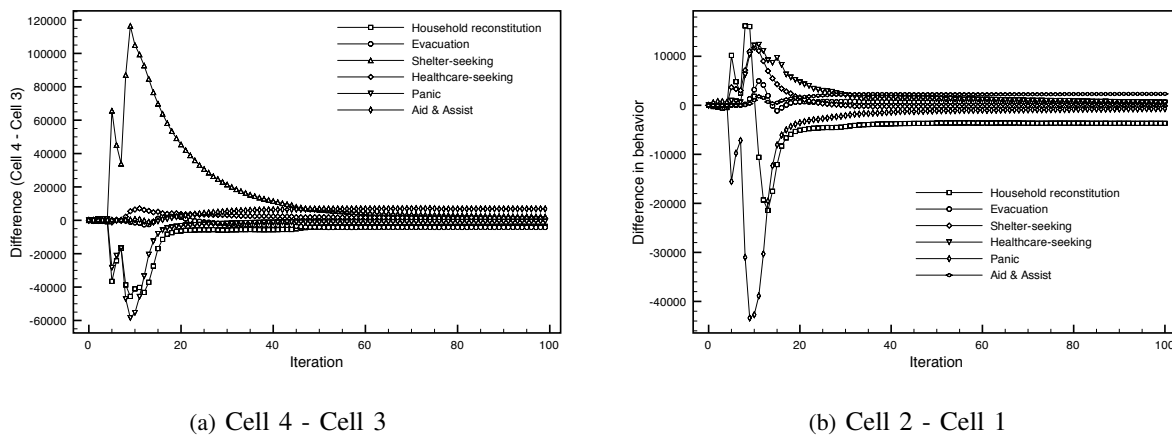


Figure 4: Differences between cells in the numbers of agents in each behavior over the course of the simulation.

Figure 4a shows the difference between cells 4 and 3 in the numbers of people engaged in each behavior over the course of the simulation. Figure 4b does the same for cells 2 and 1.

Figure 4a shows that there is a large increase early on in the number of people seeking shelter in cell 4 as compared to cell 3. There is also an early decrease in the number of people panicking or searching for family members, which are both behaviors that lead to increased exposure and risk of injury. The

combination of these behavioral changes leads to a large benefit, with more than 1500 extra people on average in good health in cell 4 at the end of two days.

Figure 4b shows that the benefit in cell 2 over cell 1 comes about for a slightly different reason. Though there is a similar drop in panic and household reconstitution (briefly), there isn't the accompanying increase in shelter-seeking because the probability of shelter-seeking in cell 2 is small. Instead we see increases in all the other behaviors. It can be shown, through a regression analysis, that the health benefit comes from more people switching to the Aid & Assist option. A well-known fact about disasters is that most of the initial search and rescue work is done by survivors and not by emergency responders. It is an affirmation of our model that this effect pops out in the simulations. More details of the comparison between cells 2 and 1 can be found in (Parikh et al. 2013).

6 VISUAL & DATA ANALYTICS

In-silico modeling and decision making that work with complex models such as the ones discussed in the preceding sections require analytical tools to interact with the system in various ways. We describe three such tools that we have developed. The first builds on a timeline environment built by researchers at MIT (simili). The extended tool allows us to specify the complex scenario. Extensions were made to the base tool so that tool can fruitfully interact with a backend database. Additional extensions included the ability to code events using scripting language, adding multiple events, etc. We built a second tool to track the progress of the models during execution. Finally, a visual analytic tool called *Dynamic Behavior Visualizer* was developed to support an analyst who wishes to view the results of the runs at the level of an individual. Snapshots of the two tools are shown in Figure 5. Together these tools significantly improve the ability of an analyst to navigate through complex scenarios and simulation outputs. The tools were valuable in detecting errors in our simulations as well.

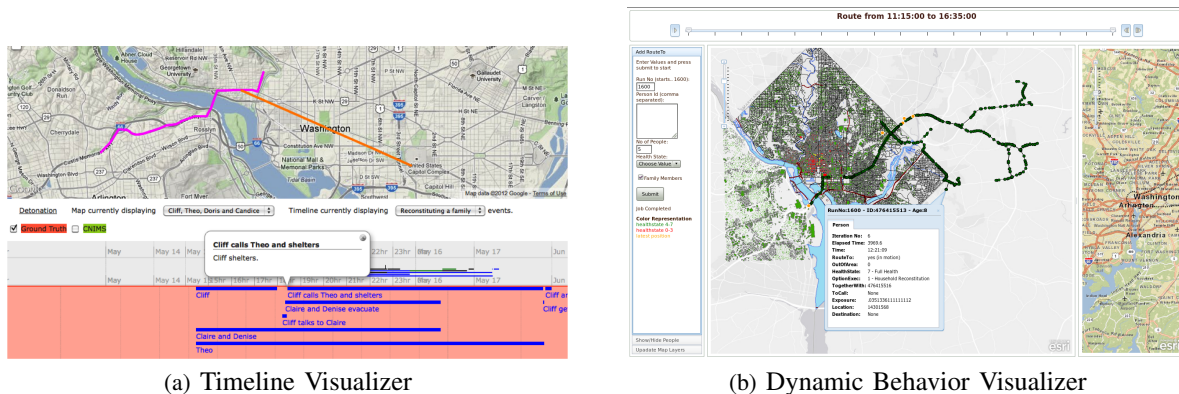


Figure 5: Snapshots of the two visual analytics tools.

7 CONCLUSION

We presented a review of a high performance computing oriented synthetic information and modeling environment that can support decision making pertaining to a large human-initiated crisis. As an illustration, we focused on a hypothetical scenario involving ground detonation caused by an improvised nuclear device in downtown Washington DC. It was based on the National Planning Scenario published by the DHS and also a subject of displaying of several prior studies. In contrast to earlier work, our focus is to understand the co-evolving individual and collective behavior in the aftermath of such a physically dominated crisis.

We outlined a scalable computing architecture that allows us to represent inter-connected urban scale civil infrastructures and the dynamic population that uses the infrastructures. A novel aspect of the

representation is its scale and fidelity. Specifically, the system can scale to represent Washington DC and its adjoining areas. Individuals and their families within the study area were represented explicitly, leading to very large co-evolving multi-networks. The environment allows us understand detailed spatial, temporal and individual level patterns of human behavior and their complex relationship with partially operational infrastructure and logistic networks.

The computational counter-factual experiments carried out using the informatics environment yield potentially interesting policy insights. They show that smart and targeted interventions can lead to substantially improved outcomes in terms of human health, affects and injuries in spite of a physically dominated event such as an IND explosion. The work reported here is just a beginning in terms of developing high resolution decision support tools for such applications.

ACKNOWLEDGMENTS

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.

REFERENCES

- Adiga, A., H. S. Mortveit, and S. Wu. 2013, May 7. "Route Stability in Large-Scale Transportation Systems". In *The Workshop on Multiagent Interaction Networks (MAIN), held in conjunction with AAMAS 2013*. St. Paul, MN, USA.
- Barrett, C., D. Beckman, M. Khan, V. S. A. Kumar, M. Marathe, P. Stretz, T. Dutta, and B. Lewis. 2009. "Generation and analysis of large synthetic social contact networks". In *Winter Simulation Conference*.
- Barrett, C., H. B. Hunt III, M. V. Marathe, S. S. Ravi, D. J. Rosenkrantz, and R. E. Stearns. 2003. "Reachability problems for sequential dynamical systems with threshold functions". *Theoretical Computer Science* 1-3:41–64.
- Barrett, C. L., R. J. Beckman, K. P. Berkgigler, K. R. Bisset, B. W. Bush, K. Campbell, S. Eubank, K. M. Henson, J. M. Hurford, D. A. Kubicek, M. V. Marathe, P. R. Romero, J. P. Smith, L. L. Smith, P. L. Speckman, P. E. Stretz, G. L. Thayer, E. V. Eeckhout, , and M. D. Williams. 2001. "TRANSIMS: Transportation Analysis Simulation System." Unclassified Technical Report No. Los Alamos National Laboratory LA-UR-00-1725, Los Alamos National Laboratory.
- Barrett, C. L., K. Bisset, S. Eubank, V. S. A. Kumar, M. V. Marathe, and H. Mortveit. 2007. "Modeling and simulation of large biological, information and socio-technical systems: An interaction-based approach". In *Proceedings of the Short Course on Modeling and Simulation of Biological Networks*.
- Barrett, C. L., M. V. Marathe, S. G. Eubank, A. K. S. Vullikanti, K. Bisset, A. Feng, M. Khan, H. Mortveit, and P. Stretz. 2012. "Social, Health and Socio-technical effects of an IND in the National Capitol: Demonstration by expanding CNIMS to involve detailed interdependent behavioral properties in National Planning Scenario 1". Technical Report 12-025, Network Dynamics and Simulation Science Lab, Virginia Bioinformatics Institute, Virginia Tech.
- Beckman, R., K. Baggerly, and M. McKay. 1996. "Creating synthetic base-line populations". *Transportation Research Part A– Policy and Practice* 30:415–429.
- Beckman, R., K. Channakeshava, F. Huang, J. Kim, A. Marathe, M. Marathe, G. Pei, S. Saha, and A. K. S. Vullikanti. 2013, June. "Integrated Multi-Network Modeling Environment for Spectrum Management". *IEEE Journal on Selected Areas in Communications: Special Issue on Network Science* 31 (6): 1158–1168.
- Buddemeier, B. R., J. E. Valentine, K. K. Millage, and L. D. Brandt. 2011, November. "National Capital Region: Key Response Planning Factors for the Aftermath of Nuclear Terrorism". Technical Report LLNL-TR-512111, Lawrence Livermore National Lab.

- Chandan, S., S. Saha, C. Barrett, S. Eubank, A. Marathe, M. Marathe, S. Swarup, and A. K. Vullikanti. 2013, April 2-5. "Modeling the Interactions between Emergency Communications and Behavior in the Aftermath of a Disaster". In *The International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction (SBP)*. Washington DC, USA.
- Chowell, G., J. M. Hyman, S. Eubank, and C. Castillo-Chavez. 2003. "Scaling laws for the movement of people between locations in a large city". *Physical Review E* 68 (066102).
- Collier, N., and M. North. 2011. "Repast HPC: A Platform for Large-scale Agent-based Modeling". In *Large-Scale Computing Techniques for Complex System Simulations*, edited by W. Dubitzky, K. Kurowski, and B. Schott, Chapter 5, 81–110. Wiley–IEEE Computer Society.
- Eubank, S., H. Guclu, V. S. A. Kumar, M. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. 2004. "Modeling disease outbreaks in realistic urban social networks". *Nature* 429 (6998): 180–184.
- Eubank, S., V. S. A. Kumar, M. Marathe, A. Srinivasan, and N. Wang. 2006. "Structure of social networks and their impact on epidemics". In *Discrete Methods in Epidemiology*, edited by G. Abello J, Cormode, Volume 70 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 179–185.
- Federal Emergency Management Agency 2010. "Planning Guidance for response to a Nuclear Detonation". *National Security Staff, Interagency Policy Subcoordination Committee for Preparedness and Response to Radiological and Nuclear Threats*.
- Gonzalez, M. C., C. A. Hidalgo, and A.-L. Barabasi. 2008. "Understanding Individual Human Mobility Patterns". *Nature* 453:779–782.
- Heber, G., and J. Gray. 2007a. "Supporting Finite Element Analysis with a Relational Database Backend, Part I: There is Life beyond Files". *Arxiv preprint cs.DB*.
- Heber, G., and J. Gray. 2007b. "Supporting Finite Element Analysis with a Relational Database Backend, Part II: Database Design and Access". *CoRR abs/cs/0701160*.
- Hybinette, M., E. Kraemer, Y. Xiong, G. Matthews, and J. Ahmed. 2006. "SASSY: A Design for Scalable Agent-Based Simulation System Using a Distributed Discrete Event Infrastructure". In *Proceedings of the 2006 Winter Simulation Conference*, 926–933.
- Kin, W., and V. Chan. 2010. "Agent-Based Simulation Tutorial–Simulation of Emergent Behavior and Differences Between Agent-Based Simulation and Discrete Event Simulation". In *Proceedings of the 2010 Winter Simulation Conference (WSC 2010)*, 135–150.
- Lasker, R. 2004. "Redefining Readiness: Terrorism Planning Through the Eyes of the Public". *Center for the Advancement of Collaborative Strategies in Health The New York Academy of Medicine*.
- Lewis, B., S. Swarup, K. Bisset, S. Eubank, M. Marathe, and C. Barrett. 2013. "A Simulation Environment for the Dynamic Evaluation of Disaster Preparedness Policies". *The Journal of Public Health Management and Practice* to appear.
- Macal, C. M., and M. J. North. 2010. "Tutorial on Agent-based Modelling and Simulation". *Journal of Simulation* 4 (3): 151–162.
- North, M. J., N. T. Collier, J. Ozik, E. R. Tatara, C. M. Macal, M. Bragen, and P. Sydelko. 2013. "Complex Adaptive Systems Modeling with Repast Symphony". *Complex Adaptive Systems Modeling* 1 (1): 3.
- Parikh, N., S. Swarup, P. E. Stretz, C. M. Rivers, B. L. Lewis, M. V. Marathe, S. G. Eubank, C. L. Barrett, K. Lum, and Y. Chungbaek. 2013, May. "Modeling Human Behavior in the Aftermath of a Hypothetical Improvised Nuclear Detonation". In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Saint Paul, MN, USA.
- Riley, P. F., and G. F. Riley. 2003. "SPADES–A Distributed Agent Simulation Environment with Software-in-the-Loop Execution". In *Proceedings of the 2003 Winter Simulation Conference (WSC 2003)*, 817–825.
- Sutton, R., D. Precup, and S. Singh. 1999. "Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning". *Artificial Intelligence* 112 (1-2): 181–211.
- Wein, L. M., Y. Choi, and S. Denuit. 2010. "Analyzing Evacuation Versus Shelter-in-Place Strategies After a Terrorist Nuclear Detonation". *Risk Analysis* 30 (9): 1315–1327.