

Resilience in the Face of Disaster: Accounting for Varying Disaster Magnitudes, Resource Topologies, and (Sub)Population Distributions in the PLAN C Emergency Planning Tool

Giuseppe Narzisi¹, Joshua S. Mincer¹, Silas Smith², and Bud Mishra¹

¹ Courant Institute of Mathematical Sciences, New York, NY 10012, USA
{narzisi, joshua.mincer, mishra}@nyu.edu

² NYU School of Medicine, NYC Poison Control Center, New York, USA
smiths11@med.nyu.edu

Project home page: <http://bioinformatics.nyu.edu/Projects/planc>

Abstract. PLAN C, an Agent-Based Model platform for urban disaster simulation and emergency planning, features a variety of reality-based agents interacting on a realistic city map and can simulate the complex dynamics of emergency responses in different urban catastrophe scenarios. Work reported here focuses on the incorporation of specific subpopulations of person agents, reflecting the existence of individuals with specific defining characteristics and needs, and their interactions with the available resources. Performance of these subpopulations are compared in both point-source attack and distributed disaster scenarios for disasters of different magnitudes. Specific “recovery points” can be derived both for total- and sub-populations, which estimate the duration of a response system’s/city’s vulnerability. The effect of varying topologies of available resources, i.e. different hospital maps, provides particular insight into the dynamics that can emerge in this complex system. PLAN C produces interesting emergent behavior which is often consistent with the literature on emergency medicine of previous events.

Key words: PLAN C, Agent-Based Modeling, Complex Systems, Disaster Management.

1 Multi-Agent and Complex Systems: Application to Disaster Management

A central problem in disaster management is the complexity inherent in an emergency response. As such, planners often rely on experience gained from previous events or drills, where possible, coupled with expert opinion. Such thinking informs policies governing operations for hospitals, responders, ambulances, etc. There exists, however, no practical way to test these policies and their effects within the global dynamic of the everyday milieu, i.e. a city with all of its resources and interacting population. The resulting gap between the theory of

disaster management and its practice is thus largely due to the fact that the real-world environment is fundamentally a *Complex System* whose intricacies cannot be reduced without sacrificing realism. Predicting the dynamics of this system, and in particular the experience of a disaster for particular subpopulations of individuals is therefore a nontrivial task.

Complex Systems are often characterized by agents capable of interacting with each other dynamically, often in non-linear and non-intuitive ways. Attempts to characterize their dynamics often results in partial differential equations that are difficult, if not impossible, to solve. One powerful technique for analyzing such complex systems is agent-based modeling (ABM) [16]. ABM has seen an exponential growth in the last few years for understanding the dynamical behavior of complex systems [16], including applications to economics [18], social science [6], biology, [5, 1] and several other real-world domains. It has been recently applied with success also in the area of Disaster Management and Preparedness (evacuation, traffic, epidemic, health-care, etc). In such models, the system is represented by a collection of autonomous decision-making entities called “agents”. A large multi-agent system can reproduce very complex dynamics even if the individual agents and their interactions follow simple rules of behavior. Emergent behaviors may be even more unpredictable and even counterintuitive when the agents are embedded and forced to interact in a real-world environment that introduces more communication channels, constraints, and behavioral rules.

RELATED WORKS. Though several mathematical and computational approaches have been proposed in recent years for public health and emergency response planning, no unified framework yet exists. These methods differ in various characteristics including: underlying modeling technology, level of details and assumptions, population size, scale and realism of the environment, etc. PLAN C has been developed in order to avoid various limitations of the following (not exhaustive) list of competing agent-based modeling and simulation tools: *InterSim* [17] is an epidemic model that uses a powerful mathematical model for modeling the agent interactions and the time course of the disease (e.g., the SEIRS model), but it lacks a realistic environment in which the simulation evolves. *EpiSim* [2], by contrast, is a highly detailed epidemic simulation system, but it is rather expensive in terms of computational time. Also, it can be difficult to collect reliable statistics from a significantly large number of simulation runs. Furthermore, it lacks an interactive user interface that could enhance its practical applicability. *DrillSim* [10] and *PedSim* [8] are two examples of evacuation models with topological constraints where the scale is restricted to evacuation plans of a floor of a building with a limited number of active agents. For example *DrillSim* was used in [9] to analyze a scenario with a total of twenty-eight evacuee agents. In [8] *PedSim* is used to suggest practical ways of minimizing the harmful consequences during evacuation of 400 people. It was found that a set of columns placed asymmetrically in front of the exit door can considerably reduce the number injuries.

1.1 PLAN C

Integral to the Large Scale Emergency Readiness (LaSER) project of NYU's Center for Catastrophe Preparedness and Response (CCPR³), is the development of a robust, agent-based model of the dynamics of urban catastrophes that permits evaluation of different biological, chemical, radiological, and other hazardous scenarios, alone or in combination, and their disaster dynamics. PLAN C⁴ [11–14] (Planning with Large Agent-Networks against Catastrophes), a modeling and simulation tool for catastrophic events, has been designed and tested in other areas of disaster preparedness. Built on top of the Java version of Repast 3.1 [15], it aims to integrate theories, algorithms, techniques, and technologies from essential preparedness disciplines in order to improve planning for and response to the public health and medical consequences of a mass casualty event. Using this robust and flexible modeling platform, public health, medical, or emergency management professionals can simulate, analyze, and predict consequences of a catastrophe, and thereby improve the capabilities of both the government and the private sectors to prepare for and respond to large-scale incidents. The agent-based modeling (ABM) paradigm provides a natural way to describe the behavior of the agents and also facilitates easier transfer of domain specific knowledge into the model.

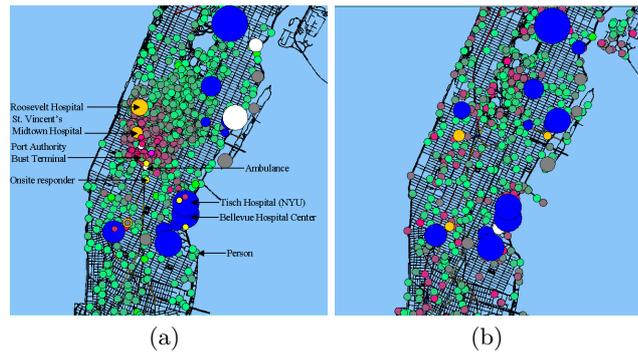


Fig. 1. PLAN C snapshots: (a) point-source scenario and (b) distributed scenario. People are labeled by small circles; color changes from green to red, signifying a high and low health level, respectively. Ambulances and first responders are labeled by small yellow circles. Larger circles signify hospitals; circle size correlates with resource level. Color represents the state of the hospital (available, critical, or full); color stays white, irrespective of state, until the first disaster patient is admitted.

PLAN C is constructed at the micro-level but permits macro-level behavior to emerge from the individual interactions of all of the agents (and their interactions, via messaging, with authorities). Behavior is not artificially imposed in

³ <http://www.nyu.edu/ccpr/>

⁴ <http://www.bioinformatics.nyu.edu/Projects/planc>

a top-down approach. In this way, the ABM technology on which PLAN C is founded is consistent with the underlying philosophy of disaster response planning: namely, planning for what people, individually and as a whole population, will do in an emergency, not what the emergency planner would like them to do [4]. Detailed descriptions of person and hospital agents, which are specifically analyzed in this research paper, are given in the appendix. A more thorough description of the PLAN C model can be found in [12, 13].

PLAN C results can be analyzed at both the macro- and micro- levels. Collective results and system dynamics at the global level (e.g. average health, fatality rate, average distress, etc.) in different emergency scenarios can be addressed. At the micro-level, analysis of selected individual agent traces allows for greater spatio-temporal resolution of disaster dynamics in a “post-debugging” process. For instance, the behavior of the hospital closest to the disaster may be compared with one further away. Similarly, the experience of a particular disaster survivor may be compared with that of a victim, perhaps enabling a finer delineation of the factors contributing to survival.

2 Simulation Results

Previous studies employed a homogeneous population of person agents. The present article reports the incorporation into PLAN C of subpopulations with special characteristics, specifically, physical disability. Dynamics of the subpopulations are analyzed and compared for a point-source attack (i.e. a sarin gas attack at the Port Authority Bus Terminal in Manhattan, NY) as well as a hypothetical distributed scenario (i.e. a previously concentrated population exposed to an agent with delayed onset of symptoms that is now distributed throughout Manhattan). Furthermore, disasters of different magnitudes vis a vis the number of casualties are considered. Finally, the effects of varying the topology of available resources (hospitals) on disaster outcome are considered for the point-source attack.

Modeling the physically disabled subpopulation is accomplished by stochastically tagging individuals with a disability factor, the value of which is initialized probabilistically from a uniform distribution between 0 and 0.5 and remains constant, reflecting the chronic nature of the disability. Thus, the range of 0 to 0.5 means that a particular disabled person can move at most half as fast as a normal person, all other factors, such as health level, being equal.

The health level (a real number in the interval $[0, 1]$) of each person in the population is initialized according to four major categories of illness defined respectively by the following probabilities and ranges: $dead_{pr} = 0.05 - [0, 0.2]$, $severe_{pr} = 0.2 - [0.2, 0.5]$, $light_{pr} = 0.3 - [0.5, 0.8]$, $no-symptoms = 0.45 - [0.1, 1]$. The incorporation of people with mild or no symptoms captures explicitly the effect of the “worried well”: people who do not actually need medical treatment but nonetheless consume available resources.

All simulations involve 20 ambulances, 5 onsite responder units and 30 hospitals. Unless otherwise stated, a total number of 10 simulations are run for each

set of initial conditions. Thus, a point on the graph represents a mean value, for that specific tick, of the parameter value found in 10 independent runs.

3 Point-source attack

The first scenario is a point-source sarin gas attack at the Port Authority Bus Terminal in Manhattan. 25% of the exposed population (either 1000, 5000, or 10000 people) suffer from impaired mobility, covering anything from low level disability (e.g., extremes of age) to chronic physical disabilities (e.g., wheelchair-bound). The emergency response and population dynamics are followed for a period of 2 days and 2 hours (3000 ticks/minutes), representing a reasonable amount of time before external aid arrives.

DIFFERENT DISASTER MAGNITUDES. Fig.2(a) features the fraction of the total population treated by each tick for disasters of different magnitudes, i.e. increasing size of the total exposed population. This parameter is an important performance measure of the emergency response system, as it characterizes its ability and efficiency in absorbing and treating disaster casualties. As seen in the figure, in the case of 1000 exposed people, 70% of the population has received treatment in the first 800 minutes of the disaster aftermath. This number decreases somewhat as the total exposed population increases. These values attest to the resilience of the city-wide emergency response system in that most victims have been seen and treated well within the first 24 hours of the disaster event. Significantly, these results compare well with actual events. The Madrid terrorist bombings of March 11, 2004 resulted in 2062 casualties. Of these, 1430 (about 70%) were treated within the first 810 minutes post event [3]. The fatality rate per tick as a percentage of the total population is graphed in the inset plot of Fig. 2(a). This value increases nonlinearly with increasing disaster magnitude.

SUBPOPULATION DYNAMICS. A finer level of detail is attained by analyzing the fraction of each subpopulation (normal and physically disabled) accessing medical treatment at each tick. As seen in Fig. 2(b), the first thing to note is that all curves are sloping downward by 3000 minutes, and in fact well before that time. All along, many people who are not very sick but nonetheless seek out medical treatment (i.e. the “worried well”) receive treatment, if necessary, and are discharged relatively quickly. This fraction shows itself in the difference between the fraction of people receiving treatment at a given tick (Fig. 2(b)) and the total fraction that has received treatment by that tick (Fig. 2(a)). There are others who require a longer stay in the hospital. As these sicker patients recover and begin to be discharged in significant numbers, the curves shown in Fig. 2(b) begin to slope downwards until all have been discharged.

Passing this point represents that the population as a whole is on the way to recovery. The timing of this recovery point is significant. A population that features an earlier point is recovering faster. In this light, the results in Fig. 2(b)

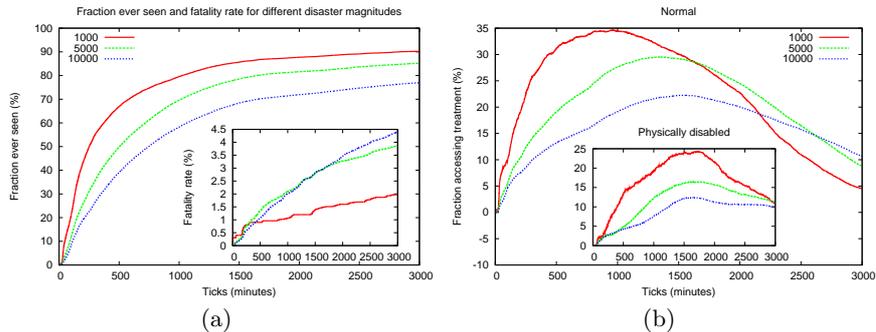


Fig. 2. Point-source attack: (a) fraction of the total population treated by each tick and fatality rate dynamics (inset plot) for the total population for different disaster magnitudes; (b) fraction of the normal and physically disabled subpopulations treated at each tick for different disaster magnitudes.

shows that the normal subpopulation as a whole recovers faster than the physically disabled population for all disaster magnitudes. Within the normal population, the point of recovery shifts to later times as disaster magnitude increases, illustrating the intuitive notion that a population recovers more slowly from a disaster of larger magnitude. Interestingly, the recovery point is fairly constant for the physically disabled subpopulation, suggesting that its ability to recover is compromised even at smaller disaster magnitudes.

DIFFERENT HOSPITAL TOPOLOGY STRUCTURES. Analysis in the aftermath of the Madrid bombings suggested that there probably existed an over-triage to the closest hospital. Namely, many noncritical patients presented to the nearest hospital [3]. In light of this conclusion, it is interesting to study the effects of the closest hospitals on disaster outcome. One way to do this is to remove sequentially hospitals close to the disaster site, i.e. to change the hospital topologies. Results in this section report on three such scenarios. In the first simulation, the closest hospital, St. Vincent's Midtown Hospital (415 W. 51st St) is removed. The second simulation features removal of this and the next closest hospital, Roosevelt Hospital (1000 Tenth Ave.). Finally, the third simulation features the additional removal of Bellevue Hospital (462 First Ave.). Fig. 1(a) shows the specific locations of these four hospitals in the map.

The fraction of each subpopulation (normal and physically disabled) accessing treatment at each tick is graphed in Fig. 3(a) for these three scenarios (and the original point-source attack with all hospitals operating). For the physically disabled population, removal of the hospitals has little effect. A relatively small, but interesting effect is seen early on for the normal subpopulation. Removal of the first hospital has very little effect, and this is consistent with the fact that St. Vincent's Midtown Hospital, while the closest, is a relatively small facility. The additional removal of Roosevelt Hospital, however, leads paradoxically to an

increase in the number of individuals receiving treatment early in the simulation and also results in an earlier recovery point. Further removal of Bellevue Hospital Center results in a slight decrease in these numbers and a slightly delayed recovery point.

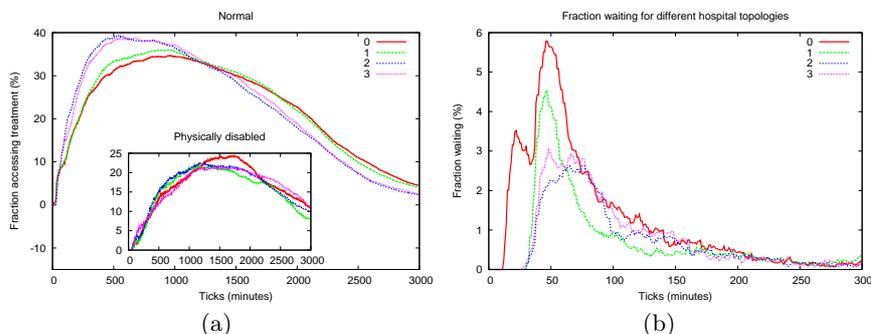


Fig. 3. Hospital topologies: (a) dynamics for the fraction accessing treatment for the normal and the physically disabled (inset plot) individual for the removal of 0, 1, 2, and 3 specific hospitals near the disaster site. (b) dynamics for the total fraction of people waiting for admission.

This trend is confirmed by analyzing the fraction of the total population that is waiting for admission to a hospital at each tick. As seen in Fig. 3(b), the number of people waiting early (first 300 minutes) in the simulation decreases significantly as the closest two hospitals are removed. Taken together, these results highlight a counterintuitive emergent phenomenon: removal of the two closest hospitals results in better performance. Further consideration suggests a reason: their removal somehow allows for a better distribution of people among the available hospitals by removing the incentive for people to move in a counterproductive fashion by crowding the closest two hospitals.

A confirmation of this hypothesis is seen in Fig. 4, which plots for each of the 30 hospitals the number of people either admitted to the hospital for treatment or waiting for admission in each minute. Each plot corresponds to the removal of 0, 1, 2, or 3 specific hospitals. As seen in Fig. 4(a), for the removal of no hospitals, crowds of people leaving the Port Authority in the aftermath of the disaster initially reach St. Vincent’s Midtown, evidenced by the early green spike in the figure. A short time later, people also arrive at Roosevelt Hospital, evidenced by the early red spike. Roosevelt Hospital, due to its large size, continues to attract a large fraction of the total population throughout the simulation. Note that the hospital curves are similar in shape to the curves for the fraction accessing treatment at each tick. They feature similar recovery points and slope downwards as the sicker people recover and are discharged. Fig. 4(b) confirms the hypothesis that dynamics do not significantly change with the removal of the closest hospital, St. Vincent’s Midtown. As seen there, the early

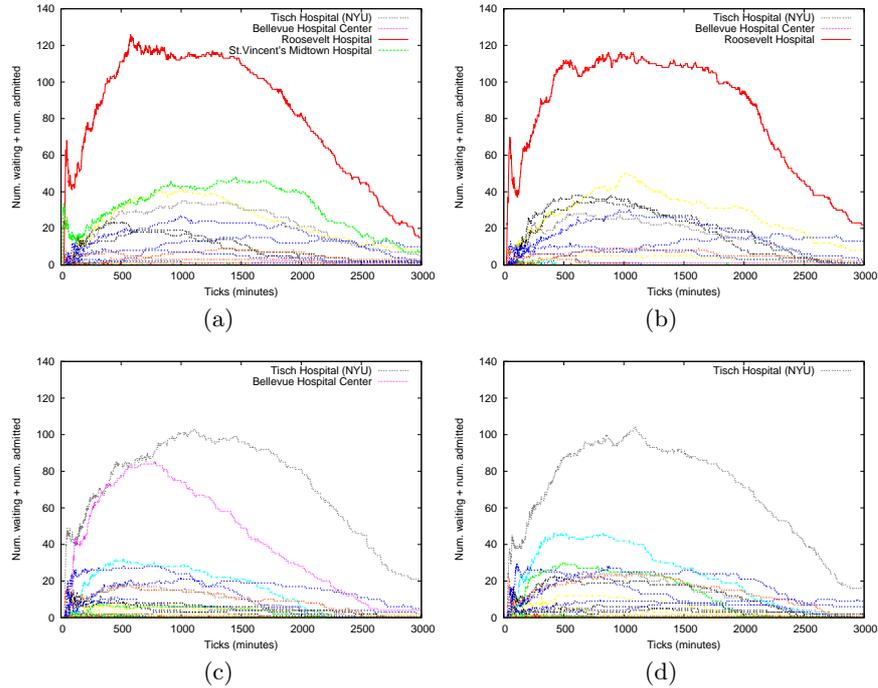


Fig. 4. Dynamics for the number of people either admitted or waiting for admission at each hospital for the case of removal of 0 (a), 1 (b), 2 (c), or 3 (d) nearby hospitals in the point-source attack scenario.

spike seen for St. Vincent's in Fig. 4(a) is no longer present, while the dynamics overall are similar to Fig. 4(a), with Roosevelt Hospital playing a dominant role. As seen in Figs. 4(c) and 4(d), people's movements change dramatically once Roosevelt Hospital is also removed. Instead of moving northward, people move east to Tisch Hospital (gray curve) and nearby Bellevue Hospital (pink curve), which is no longer seen in Fig. 4(d) as it has also been removed. Moving east instead of north results in the better outcome seen in Figs. 3(a) and 3(b). This is understood in light of the greater number of resources that are found there. For example, Tisch and Bellevue are both large hospitals in close proximity to one another, as compared with Roosevelt and St. Vincent's Midtown, whose combined resources are less. Furthermore, movement east places people closer to larger hospitals on Manhattan's East Side.

These results are particularly illustrative of the usefulness of PLAN C in emergency planning. While suggesting that the distribution of hospitals in Manhattan as given may not be optimal for a specific disaster scenario, they also suggest that emergency managers, through risk communication or other means, can guide people's movements in particular directions and to particular hospitals to improve outcome. Moreover, the map for a given city may not necessarily

be a given constraint. Different configurations for the resource locations can be designed and PLAN C could be employed to determine the optimal layout for a given scenario, enabling planners to redesign an entire city or a particular locale for optimal robustness in the face of a specific disaster.

4 Distributed scenario

In contrast to the point-source attack studied above, the current section focuses on a distributed scenario in which the individuals are positioned randomly throughout Manhattan. The hypothetical situation is one in which a large number of people were congregated at some earlier point in time. All were exposed to a hypothetical agent that does not reveal itself symptomatically until a later time. As the simulation starts, a certain number are already dead, and the city-wide health system is becoming aware of the nature of the situation. The parameters for the hypothetical agent vis a vis health decline are made identical to those of the sarin scenario in order to facilitate comparisons between a point-source and distributed attack. The distributed scenario does not feature onsite responders or ambulances, as PLAN C does not currently include ambulances that can respond to particular emergency calls.

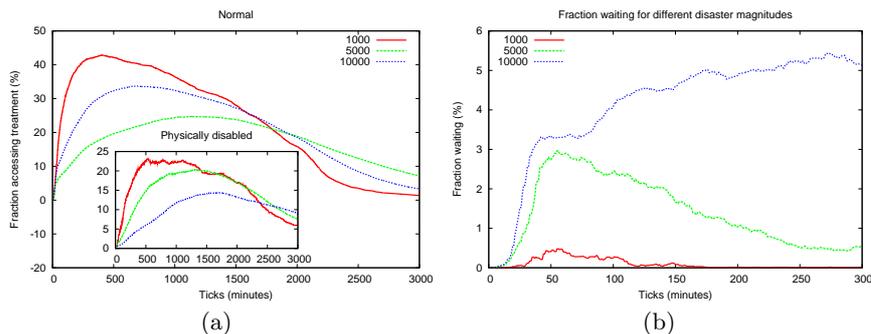


Fig. 5. Distributed scenario: (a) fraction of the normal and physically disabled subpopulations treated at each tick for different disaster magnitudes; (b) dynamics for the total fraction of people waiting for admission.

Fig. 5(a) highlights the fraction of the normal and physically disabled subpopulations that are receiving medical treatment at each tick for disasters of different magnitudes. Comparison with the point-source attack (Fig. 2) demonstrates a robustness of the physically disabled population not seen previously. Namely, the recovery point (maximum) is earlier for disasters of lesser magnitude, whereas in the point-source attack it was constant. This robustness in the distributed scenario is expected. There is less competition for the same hospital resources as people distribute more equitably to a variety of hospitals. This more

equitable distribution is seen when comparing Fig. 6, which plots the number of people admitted or waiting for admission at each hospital, with the analogous figure for the point-source attack (Fig. 4(a)). From the point of view of the city-wide health system, the ability to absorb disaster casualties is significantly compromised in a point-source attack as compared with a distributed scenario.

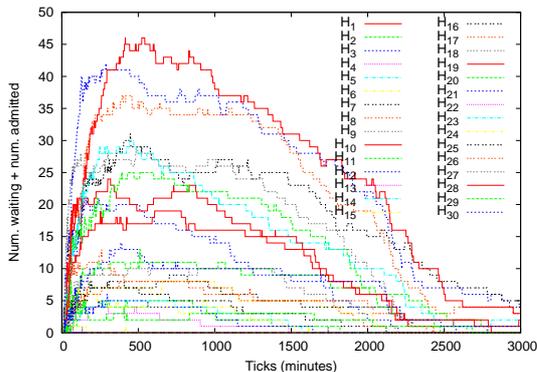


Fig. 6. Distributed scenario: dynamics for the number of people either admitted or waiting for admission at each hospital (represented by a different trace).

Analysis of the total fraction of people waiting for admission reveals an interesting result (Fig. 5(b)). For a distributed disaster magnitude of 1000 people, the number waiting increases early in the simulation and then decays. A similar form is seen for 5000, with a slower decay. Finally, for 10000 people, the curve does not decay in the early part of the simulation. Even for the distributed simulation, which yields a more robust city-wide disaster response with earlier recovery points as compared with the point-source attack, a large disaster can still stress the system.

5 Conclusions and future investigations

Results presented here illustrate the feasibility of incorporating special subpopulations within the PLAN C framework, capable of simulating disasters of various magnitudes (10000 casualties). Analysis of the system dynamics in disasters of varying topologies, both in terms of the spatial distribution of subpopulations (point-source versus distributed scenario) as well as the locations of the available resources (hospitals), can reveal counterintuitive emergent phenomena, such as the fact that removing hospitals close to the disaster site can improve overall outcome of the system. A significant finding is that “recovery points” can be discerned, both for total- and sub-populations. When carried forward with disasters of larger magnitude or of longer duration (e.g. infectious diseases), these

points can be determined in order to estimate the duration of a response system's/city's vulnerability. Proactive use of PLAN C in planning would attempt to maximize system resilience and earlier recovery points by optimizing various facets of disaster response. Emergency managers, urban planners and public health officials can refine existing emergency plans and policies using PLAN C's versatile, innovative platform.

Among the many lines of investigation which we plan to address, particularly relevant are the following: 1) greater realism can be incorporated by the introduction of social networks linking individuals across subpopulations with consequent dynamics that may prove very different as normal individuals alter their dynamics to aid "friends" or family members who suffer from some impairment such as physical disability; 2) validation of PLAN C's realism in order to build confidence in its use as a tool for disaster planning; 3) automatic computation of optimal configurations (locations) for the available resource (i.e. hospitals) through the use of multi-objective evolutionary algorithms [14].

Acknowledgments. We would like to thank Lewis Nelson, Dianne Rekow and Ian Portelli for their contribution to the clinical aspect of the study design and development. We also thank Venkatesh Mysore, Ofer Gill, Jee Woong Byeon and Raoul-Sam Daruwala who contributed to the implementation in Repast. The work was supported by New York University's Center for Catastrophe Preparedness and Response, through its U.S. Department of Homeland Security grant #2004-GT-TX-0001.

References

1. Chaitanya A. Athale and Thomas S. Deisboeck. The effects of egf-receptor density on multiscale tumor growth patterns. *Journal of Theoretical Biology*, 238(4):771–779, 2006.
2. Barrett CL, Eubank SG, and Smith JP. If smallpox strikes portland... *Scientific American*, 292(3):42–9, 2005.
3. Gutierrez de Ceballos JP, Turegano Fuentes F, Perez Diaz D, Sanz Sanchez M, Martin Llorente C, and Guerrero Sanz JE. Casualties treated at the closest hospital in the madrid, march 11, terrorist bombings. *Crit Care Med.*, 33(1):S107–12, 2005.
4. E. Auf der Heide. The importance of evidence-based disaster planning. *Annals of Emergency Medicine*, 47(1):34–49, 2006.
5. Thierry Emonet, Charles M. Macal, Michael J. North, Charles E. Wickersham, and Philippe Cluzel. AgentCell: a digital single-cell assay for bacterial chemotaxis. *Bioinformatics*, 21(11):2714–2721, 2005.
6. Joshua M. Epstein. *Generative Social Science: Studies in Agent-Based Computational Modeling*, volume 2. Princeton University Press, January 2007.
7. Niels K. Rathlev et al. Time series analysis of variables associated with daily mean emergency department length of stay. *Annals of Emergency Medicine*, 49(3):265–271, 2007.
8. Dirk Helbing, Illes Farkas, and Tamas Vicsek. Simulating dynamical features of escape panic. *Nature*, 407:487, 2000.

9. Daniel Massaguer, Vidhya Balasubramanian, Sharad Mehrotra, and Nalini Venkatasubramanian. Multi-agent simulation of disaster response. In *Proceedings of the First International Workshop on Agent Technology for Disaster Management, ATDM*, 2006.
10. Daniel Massaguer, Vidhya Balasubramanian, Sharad Mehrotra, and Nalini Venkatasubramanian. Synthetic humans in emergency response drills. *Demo paper in AAMAS*, 2006.
11. Venkatesh Mysore, Ofer Gill, Raoul-Sam Daruwala, Marco Antoniotti, Vijay Saraswat, and Bud Mishra. Multi-agent modeling and analysis of the brazilian food-poisoning scenario. In *The Agent Conference*, 2005.
12. Venkatesh Mysore, Giuseppe Narzisi, Lewis Nelson, Dianne Rekow, Marc Triola, Alan Shapiro, Clare Coleman, Ofer Gill, Raoul-Sam Daruwala, and Bud Mishra. Agent modeling of a sarin attack in manhattan. In *Proceedings of the First International Workshop on Agent Technology for Disaster Management, ATDM*, 2006.
13. G. Narzisi, V. Mysore, L. Nelson, D. Rekow, M. Triola, L. Halcomb, I. Portelli, and B. Mishra. Complexities, catastrophes and cities: Unraveling emergency dynamics. *International Conference on Complex Systems (ICCS)*, 2006.
14. Giuseppe Narzisi, Venkatesh Mysore, and Bud Mishra. Multi-objective evolutionary optimization of agent based models: an application to emergency response planning. In *Computational Intelligence - CI 2006*, pages 224–230. ACTA Press, 2006.
15. M.J. North, N.T. Collier, and J.R. Vos. Experiences creating three implementations of the repast agent modeling toolkit. *ACM Transactions on Modeling and Computer Simulation*, 16(1):1–25, 2006.
16. PNAS. *Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling*, volume 99(3). May 2002.
17. Markus Schwehm, Chris Leary, Hans-Peter Duerr, and Martin Eichner. InterSim: A network-based outbreak investigation and intervention planning tool. *International Conference on Complex Systems (ICCS)*, 2006.
18. Leigh Tesfatsion and Kenneth Judd. *Handbook of Computational Economics, Agent-Based Computational Economics*, volume 2. Elsevier, July 2006.

Appendix: PLAN C’s agent behavior description

PLAN C is an innovative agent-based framework for simulating large scale disasters in urban settings which features: (i) large number of computational actors/agents (Person, Hospital, On-Site Responder, Ambulance and Catastrophe); (ii) flexible number of parameters for describing the agents’ behavior and interaction, the time course of the disease, environmental factors, etc.; (iii) communication channels to exchange information (e.g., health/resource levels, hospital operation mode, etc.); among similar and differing agents; (iv) realistic models of medical/responder units and catastrophe chemical agent effects; and (v) integrated urban topologies (streets, subways, hospitals, etc), via publicly available GIS (geographical information system) data. PLAN C involved collaborative participation from a multi-disciplinary team including medical, sociological and legal experts from the NYU’s CCPR (Center for Catastrophe Preparedness and Response). Because of the focus of the present paper, we present detailed descriptions of the Person and the Hospital agents.

Hospital agent

The hospital is a stationary agent that is an abstraction of any medical facility that can play a role at the time of a catastrophe. The hospital uses a simple triage policy with three operation modes (*available*, *critical* and *full*). The transition from one state to another is based on the available number of beds and resources. The hospital is realistically modeled to include an Emergency Department (ED), inpatient beds, isolation beds, and critical care capacity in the form of an Intensive Care Unit (ICU) and ventilators.

The hospital mode directly influences several decisions (*triage*): whom to turn away, whom to admit, whom to treat, how much resources to allocate to a person requiring treatment, who can be discharged, and who can be moved from the ED to an inpatient or ICU bed. In the available mode, the hospital admits all persons present for treatment; in the critical mode, only critically ill people can be admitted; in the full mode, no new people can be admitted. Among the several parameters which influence the way the hospital operates, the following are particularly relevant. Thirty major hospitals have been included, and the number of hospital beds was used as an indicator of the capacity of the hospital.

PROBABILITY OF ADMISSION. This equation captures the assumption that increased hospital efficiency results in a higher probability of admitting additional people in each tick. Conversely, as the number of occupied beds increases, the admission probability decreases:

$$P_r(\text{admit}) = \frac{Eff_H \times rate_p \times tick_{size}}{1 + |H|_{occ}} \quad (1)$$

where Eff_H is the efficiency of the hospital H (defined later), $tick_{size}$ is the number of minutes per tick of the simulation ($= 1$), $rate_p$ ($= 50$) is the number of people per minute that can be attended to (admitted + treated + discharged), and $|H|_{occ}$ is the number of total occupied beds in the hospital.

HOSPITAL EFFICIENCY. The efficiency, regulated by the following rule, directly influences many decisions, such as the amount of treatment given to a person as well as the number of persons admitted or treated at each tick, also it indirectly affects the waiting time at the hospital and mean hospitalization time:

$$Eff_H = \frac{1}{1 + \left(\frac{|H|_{occ} + S_H}{|H| + |ED|} \right)} \quad (2)$$

$|H|_{occ}$ is defined as before, while $|H|$ and $|ED|$ are respectively the total number of inpatient and ED beds and S_H is the level of *sickness* inside the hospital, defined as $S_H = \sum_{i:p \in H} (1 - h_p)$, where h_p is the health level of the person p . The intuitive notion behind this equation is that the efficiency of the hospital should decrease as bed occupancy and the overall sickness of the inpatient population increases. We are currently studying different variations of these formulas to better model these processes.

AVERAGE TIME IN ED. The average time a person spends in th ED before being moved to an inpatient bed is defined by the following equation:

$$Avg_t^{ED} = base_t \times (1 + |Ed|_{occ} + \frac{S_{ED}}{|ED|}) \quad (3)$$

where $base_t$ ($= 5$) is a user parameter defining the base value for the average time in the ED, S_{ED} is the total sickness of people in the ED. This formula and the previous one are consistent with the time series analysis in [7].

Person agent

The affected population is modeled as reactive *selfish* agents with *bounded rationality* and *stochastic behavior*. The person's initial goal is to reach the original destination (home or place of work) from the initial location. After the event, health begins to deteriorate such that at a certain health-level, governed by environmental and personality factors, the person decides to head to a hospital. The person agent maintains information about destination (home/work or hospital), current health level $h_l \in [0, 1]$, current level of medical intervention $m_l \in [0, 1]$; location and current capacity of known hospitals. An agent may talk to any agent in its neighborhood (defined as everybody in the same location on the map), and exchange information about the known list of hospitals, the disease type, etc.

LEVEL OF WORRY AND COMPLIANCE. Each person has specific personality traits defined through the degree of worry/fear $w_l \in [0, 1]$, which represents the innate level of irrationality in the agent, and the level of compliance $c_l \in [0, 1]$, which captures the instruction-abiding trait of a person; Both w_l and c_l are initialized uniformly random in $[0, 1]$ but they also change during the simulation as a consequence of the interactions between the agents: when two agents talk to each other they both update the variables w_l and c_l according to the current values of the other agent computing the mean value.

LEVEL OF DISTRESS. The degree of worry and the health level are combined together to define the perceived level of distress of a person: $s_l = w_l \times (1 - h_l)$. The simple intuition behind this formula is the following: if the health level is high then with low probability the degree of worry can generate distress. This parameter influence many decisions of the person agent, for example, higher the level of distress suffered by a person, higher the probability of selecting the nearest hospital even when it is full.

DISABILITY FACTOR. The disability factor d_l reflects the chronic nature of the disability of the person. This parameter which is initialized randomly in $[0, l]$, where l is used to decide the degree of disability, is then used as a multiplication factor with other characteristics parameter of a person. In this paper, the speed of a disabled person is updated proportionally to d_l .