Chapter 1

Complexities, Catastrophes and Cities: Unraveling Emergency Dynamics

Giuseppe Narzisi
University of Catania, narzisi@dmi.unict.it
Courant Institute, NYU, narzisi@nyu.edu
Venkatesh Mysore
Courant Institute, NYU, mysore@cs.nyu.edu
Lewis Nelson
NYU School of Medicine, lnelson@pol.net
Dianne Rekow
NYU College of Dentistry, edr1@nyu.edu
Marc Triola
NYU School of Medicine, marc.triola@med.nyu.edu
Liza Halcomb
NYU School of Medicine, halcoe01@med.nyu.edu
Ian Portelli
NYU CCPR, ian.portelli@nyu.edu
Bud Mishra
NYU, 715 Broadway #1002, New York, NY, USA, mishra@nyu.edu

[1]

[1]Author contributions: B.M. designed research; G.N., V.M. and B.M. performed research; G.N., V.M. and B.M. contributed new analytical tools; LN, DR and MT contributed to the clinical aspects of the study design and development; G.N., V.M. and B.M. wrote the paper; and G.N., V.M., L.N., D.R., M.T., L.H., and I.P. reviewed the paper.

1 Introduction & Background

Complex Systems are often characterized by *agents* capable of interacting with each other dynamically, often in non-linear and non-intuitive ways. Trying to characterize their dynamics often results in partial differential equations that are difficult, if not impossible, to solve. A large city or a city-state is an example of such an evolving and self-organizing complex environment that efficiently adapts to different and numerous incremental changes to its social, cultural and technological infrastructure [2]. One powerful technique for analyzing such complex systems is Agent-Based Modeling (ABM) [10], which has seen an increasing number of applications in social science, economics and also biology. The agent-based paradigm facilitates easier transfer of domain specific knowledge into a model. ABM provides a *natural* way to describe systems in which the overall dynamics can be described as the result of the behavior of populations of autonomous components: agents, with a fixed set of rules based on local information and possible central control. As part of the NYU Center for Catastrophe Preparedness and Response (CCPR[2]), we have been exploring how ABM can serve as a powerful simulation technique for analyzing large-scale urban disasters. The central problem in Disaster Management is that it is not immediately apparent whether the current urban emergency plans are robust against such sudden, rare and punctuated catastrophic events. An agent-based emergency response model can utilize the large amount of information about the possible rules of behavior for people, hospitals, on-site responders and ambulances, without depending on the scarce knowledge about the efficacy of those rules or the global dynamics.

We have been striving towards a methodical and algorithmic approach for both preparedness and response, by combining powerful ideas from model-checking, simulation and multi-objective optimization, in order that a large urban structure can recover from the effects of a disastrous event quickly and efficiently. More recently, game theoretic paradigms have also influenced the analysis of complex systems. In our models, persons play “games” with each other for the medical resources; persons and hospitals interact to minimize several factors like number of fatalities, average waiting time, average ill-health, cost, etc. Likewise, the heuristics people employ to choose the hospital they should head to, based on prior knowledge about their size and location and real-time knowledge about current occupancies, can be seen as an extension of the Santa Fe bar problem [4]. Game theory also discusses different kinds of strategies that can effectively describe different personality, cultural and social traits governing panic behavior: some people just imitate their neighbors, some are contrarian, some are rational, some are irrational, some employ a random strategy, etc.

Disaster planning is often based on assumptions derived from a conventional wisdom that is at variance with empirical field disaster research studies [3]. Our

[2]http://www.nyu.edu/ccpr

efforts to avert this error have resulted in a new system[3] with well-identified, validated, simple rules with minimal number of parameters to avoid modeler bias and unnecessary complexity. The persons, hospitals, on-site responders, ambulances and disease prognosis follow deterministic rules with probabilistic parameters that can be modified by the user. A more detailed description of our system can be found in [8], where the Sarin gas exposure scenario is investigated in the constraints defined by Manhattan, New York, and in [7], where the Brazilian food poisoning scenario is recreated. The system is implemented in Repast 3.1 [4] [9], a popular and versatile Java-Based software toolkit that has been used to model such diverse concepts like intracellular processes and business strategies. We have also integrated ProActive [5] with RePast, in order to use the computational power of a cluster of computers to explore the parameter space of the system. Rather than focusing on the intricacies of the modeling problem, in this paper, we delve into the nature and sources of complexity in the dynamics of different kinds of catastrophes.

2 Experimental results

In disaster management, it has been established that "Planning should take into consideration how people and organizations are likely to act, rather than expecting them to change their behavior to conform to the plan" [3]. ABM serves as a means of describing the behavior of medical facilities (controllable) and evaluating their performance in different disease scenarios for people with different personality and health profiles. Unless stated otherwise, experimental results are carried out using the same values for the parameters as described in table 1 and each plot is averaged on 10 independent runs.

Single event scenario As a first scenario, we consider a possible terrorist attack with a warfare agent at Port Authority Bus Terminal in the island of Manhattan. In order to understand the complexity of the system dynamics, in Fig. 1, we monitor different statistics for the affected population. The left plot in Fig. 1 shows the evolution curves for the average waiting time of the affected population at the hospitals. The presence of three jumps is visible in the first 400 ticks of the curves, corresponding to the crowding effect of the flux of people at the three nearest hospitals to the site of the attack. Each climb phase is a consequence of the hospital state changing rapidly from "available" to "critical," with a resulting increase in the number of waiting non-critical persons. The flat phase that ensues is due to the state change from "critical" to "full", where all waiting persons are instructed to head to another hospital. It is interesting to note how the population size of 500 persons seems to produce a more complex scenario as compared to a large size of 1000, as evident in the higher waiting time

[3] The details of the system have been summarized in the *Appendix*.
[4] http://repast.sourceforge.net
[5] http://www-sop.inria.fr/oasis/proactive/

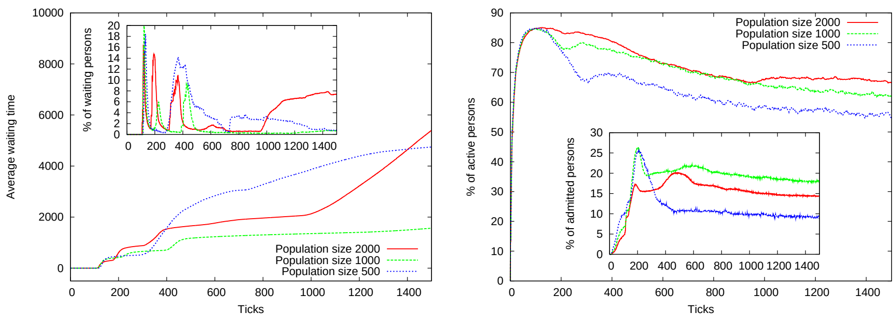

Figure 1: Left plot: evolution curves for the percentage of waiting persons at the hospitals and the average waiting time of the population with population size 500, 1000 and 2000. Right plot: evolution curves for the percentage of active and admitted persons with population size 500, 1000 and 2000.

at the hospitals. This unforeseen outcome can be explained by observing that after the nearest hospital becomes full, the remaining waiting population that heads to another hospital is unable to fill up the new one. The new hospital remains in a critical state for more time causing an increased waiting time. This effect is visible in the inset plot on the left of Fig. 1, where the curve for the population size of 500 produces the highest percentage of waiting persons around 400 ticks. A similar behavior is produced by an affected population of 2000 individuals, but in this case the scenario unfolded after the three nearest hospitals became full.

The right plot of Fig. 1 shows the percentage of active and admitted persons. The term *active* denotes a person who has decided to head to a hospital. As expected, immediately after the attack, both the number of active and admitted persons quickly increases, but then different courses are produced by the different population sizes. Another unexpected behavior emerges in the right inset plot of Fig. 1: an affected population of 1000 individuals produces a higher percentage of admitted persons than that of 2000. A possible explanation can be found by observing that the resources of each hospital are the same for both population sizes, but the number of persons with lethal and severe injuries increases with the population size. These are persons who need more treatment producing a longer hospitalization time and higher demand of resources. At the same time, there are also many persons, some lightly and others severely injured, who are awaiting admission.

Multiple event scenario As a second scenario, we consider a possible terrorist attack involving multiple explosions – in particular, caused by three bombs located respectively in Union Square, Times Square and Central Park. The explosions are simulated to occur after 10, 120 and 300 minutes respectively. A population of 5000 persons is involved and initialized to random positions on the map at the beginning of the simulation. The left plot of Fig. 2 shows the

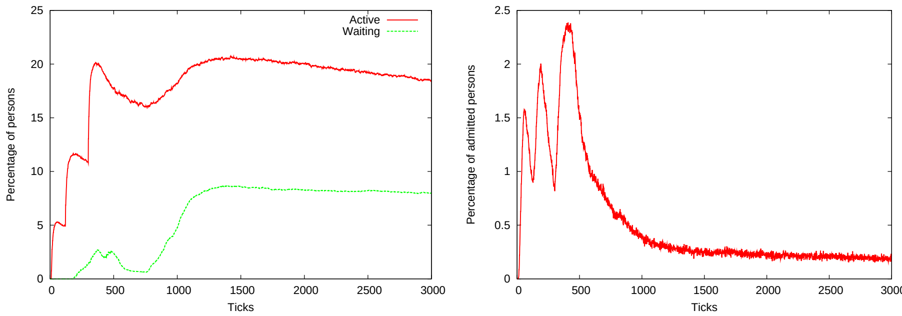

Figure 2: Left plot: evolution curves for active and waiting persons. Right plot: evolution curve for the percentage of admitted persons in the hospitals.

expected increase in the number of active persons after each of the three explosions. It is interesting to note the presence of an unpredictable fourth but less rapid increase after 1000 ticks. The waiting curve instead follows a completely different path because of the different spatial positions of the hospitals with respect to the sites of the explosions and their different resource levels. The right plot of Fig. 2 shows the curve for the percentage of admitted persons in the hospitals. As expected, after each explosion we have an increase in the number of admissions, but most of them are probably persons who do not need long-term hospitalization and hence, are discharged soon. However, the percentage of admitted persons never becomes zero; random fluctuations after the 700^{th} tick are visible due to the probabilistic personality factors (irrationality and obedience) of each person.

2.1 ABM model-checking

Unlike statistical analysis of metrics averaged over multiple agents and simulations, the model-checking approach focuses on individual agents' traces. Complex temporal properties may be described in Linear Temporal Logic (LTL) and then model-checked in a model-checker such as XSSYS[1]. The trace analysis can help discover finer aspects about the underlying system dynamics. The XSSYS system was originally developed for simulating and analyzing biochemical pathways. The agents' traces produced in output by the system can be read using XSSYS. To demonstrate the technique, we consider an intensive toxic agent exposure in downtown Manhattan and monitor a person and a hospital in Fig. 3.

3 Conclusions and future investigations

The complex interactions between the affected population and the available resources of a response plan have remained poorly understood, are still beyond the analytical capability of traditional modeling tools, and have resisted any system-

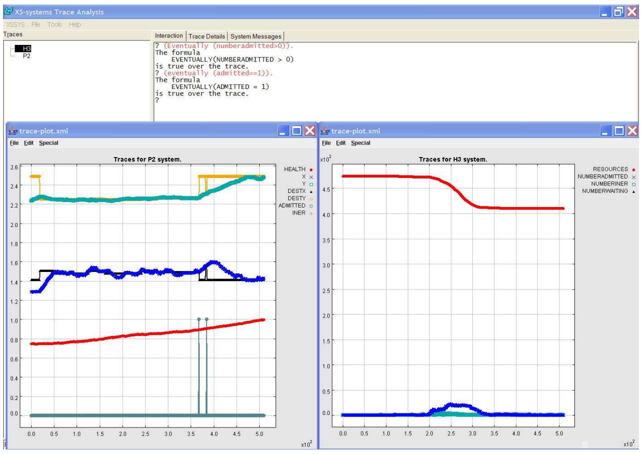

Figure 3: Temporal Logic Analysis in XSSYS. Left plot: Time-Trace of a Person. Right plot: Time-Trace of a Hospital.

atic investigation. In this research work we have shown that a deep analysis of the source of complexity generated by the simulation of different kind of urban emergency scenarios is effectively possible. This finer analysis has been accomplished using large-scale simulation with a novel Agent-Based Model simulation tool[6] addressing disasters in urban settings. Simple rules of behavior are seen to produce uncanny emergent dynamics with unpredictable interdependences, which, with the help of the statistical analysis features of the system, can be inspected in order to refine existing plans and policies.

Currently, we are extending the model in order to simulate not just immediate one-time events, e.g., attack with a chemical agent, bomb explosion, etc., but also long-lasting slowly-unfolding scenarios such as those resulting from an infectious disease, e.g., Smallpox. We are also working on integrating the Repast toolkit with the rule engine JESS[7]. This will facilitate characterization of agent behavior with declarative rules and run-time modification. Moreover, response plans involve different, often conflicting, criteria that must be satisfied and optimized in parallel – number of fatalities, average population health, time taken to succumb, waiting time at the hospital, etc. In our framework, a response plan is expressed in terms of the system rules and parameters, producing a gargantuan strategy space that should be explored in order to find "optimal" plans. We are currently exploring the use of *multi-objective evolutionary algorithms* to address this computational challenge.

Through our efforts, we want to demonstrate that the ABM paradigm, in conjunction with statistical analysis, multi-objective optimization, game theory and model-checking of agent-traces, offers a novel way to understand, plan and control the unwieldy dynamics of a large-scale urban emergency response.

[6] http://www.bioinformatics.nyu.edu/Projects/planc/index.shtml
[7] http://herzberg.ca.sandia.gov/jess/

Acknowledgments

We would like to acknowledge the funding support from Department of Homeland Security Grant # 2204-GT-TX-0001. Ofer Gill and Raoul-Sam Daruwala also contributed to the implementation in RePast.

Bibliography

[1] Antoniotti, M., F. Park, A. Policriti, N. Ugel, and B. Mishra, "Foundations of a query and simulation system for the modeling of biochemical and biological processes", *The Pacific Symposium on Biocomputing: PSB* (R. Altman, A. Dunker, L. Hunter, T. Jung, and T. Klein eds.), World Scientific (January 2003), 116–127.

[2] Batty, Michael, *Cities and Complexity: Understanding Cities with Cellular Automata, Agent-Based Models, and Fractals*, MIT Press (2005).

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

[4] Greenwald, A., B. Mishra, and R. Parikh, "The santa fe bar problem revisited: Theoretical and practical implications", *The Proceedings of the Summer Festival on Game Theory: International Conference*, (1998).

[5] Hick, John L., and Daniel T. O'Laughlin, "Concept of Operations for Triage of Mechanical Ventilation in an Epidemic", *Acad Emerg Med* (2006), j.aem.2005.07.037.

[6] Korf, R.E., "Real-time heuristic search", *Artificial Intelligence* **42** (1990), 189–211.

[7] Mysore, Venkatesh, Ofer Gill, Raoul-Sam Daruwala, Marco Antoniotti, Vijay Saraswat, and Bud Mishra, "Multi-agent modeling and analysis of the brazilian food-poisoning scenario", *The Agent Conference*, (2005).

[8] Mysore, Venkatesh, 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).

[9] North M.J., Collier N.T., and Vos J.R., "Experiences creating three implementations of the repast agent modeling toolkit", *ACM Transactions on Modeling and Computer Simulation* **16**, 1 (2006), 1–25.

[10] PNAS, *Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling* vol. 99(3), (May 2002).

Appendix

One of the main issues in ABM is to build models at the appropriate level of description, using the requisite level of details in order to produce a system that serves its analytical purpose. The details of our model have been summarized below from [8]. The table 1 shows the main parameters that the user can modify.

Table 1: Main model parameters.

Name	Range	value
Maximum number of iterations	$[0, \infty]$	2000
Number of agents (Person, Hospitals, On-site Responders, Ambulances)	$[0, \infty]$	(1000, 28, 5, 10)
Alert time (in minutes, for On-site responders and Ambulances)	$[0, \infty]$	30
Critical health Level	[0, 1]	0.3
Non-critical health level	[0, 1]	0.6
Unsafe health level	[0, 1]	0.4
On-site Responder dischargeable health level	[0, 1]	0.5
Hospital dischargeable health level	[0, 1]	0.8
Probability to have a communication device (phone, radio, etc)	[0, 1]	1.0
Phone update probability	[0, 1]	0.2
Hospital low resource level (percentage of)	[0, 100]	30
Hospital very low resource level (percentage of)	[0, 100]	10
Hospital low beds level (percentage of)	[0, 100]	30
Probability of lethal, severe, light, injuries	[0, 100]	0.1, 0.35, 0.4, 0.15

PERSON: 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 catastrophic event, his/her health begins to deteriorate and at a certain health-level, decided by environmental and personality factors, the person decides to head to a hospital. The person agent maintains information about: (1) the destination (home/work or hospital); (2) current health level $h_l \in [0, 1]$; (3) current level of medication $t_l \in [0, 1]$; (4) location and current capacity of known hospital; (5) degree of worry $w_l \in [0, 1]$ that represents the innate level of irrationality in the agent; (6) level of obedience $o_l \in [0, 1]$ that captures the instruction-abiding trait of a person; and (7) perceived level of distress $d = w_l \times (1 - h_l)$. Higher the level of distress suffered by a person, higher would be the probability of selecting the nearest hospital even when it is full; similarly, higher the level of obedience in a person, higher would be the probability of following the rules (heading to a hospital at the correct health level specified by the user).

HOSPITAL: It is modeled as a stationary agent that is an abstraction of any medical facility that can play a role at the time of a catastrophe. Twenty eight major hospitals have been included, and the number of hospital beds was used as an indicator of the capacity of the hospital. The hospital operates in three modes: *available*, *critical* and *full* according to the residual number of beds and resources. The hospital mode directly influences several decisions: whom to turn away, whom to treat and how much resources to allocate to a person requiring treatment. The medical term for this process is *triage* [5]. In the "available" mode, the hospital admits all persons that request treatment; in the "critical" mode, only critically ill persons are admitted; in the "full" mode, nobody new is admitted.

On-Site responders: On-site treatment is provided by Major Emergency Response Vehicles (MERVs), Hazardous Materials team (HazMat) and Emergency Medical Services (EMS) ambulances. These small mobile hospitals are initially inactive and stationary at their hospital of affiliation. After receiving notification of the disaster, they move towards the catastrophe site. Their knowledge includes: (1) starting location; (2) time of dispatch; (2) locations and current capacities of known hospitals. The behavior is exactly the same as a hospital in critical mode.

Ambulance: It is modeled as a small mobile hospital that transports sick persons to the nearest hospital. The ambulance is initially at a random point in the map and does not help anybody. After being activated, their destination node is assigned to the location of the catastrophe site. After reaching the site and collecting one person, the ambulance assists in moving them to a hospital, also providing treatment during the travel.

Catastrophe: The catastrophe itself is modeled as an agent in the system. This particular implementation gives us the ability to model very different scenarios. In particular the catastrophe-agent can be specialized in order to model a source of poisoning, a bomb, a warfare agent, etc. More importantly it is possible to initialize multiple catastrophe-agent and setting their time of activation.

Disease: The time-course of the person's health after the disaster is modeled as a three step probabilistic function such that the healthier the person, the more likely that his / her health will improve rather than deteriorate:

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

where $U(0,1)$ is a real random number generated uniformly in the range $(0,1)$.

Topology and Transportation: Publicly available Geographic Information Systems (GIS) data about the roads of the city was converted into a graph, where nodes are intersections and edges are streets. We have performed this conversion for Manhattan island in New York city. Agents are constrained to move only along the edges of the graph, with the effective speed at each time-step depending on the health level and probabilistic terms to simulate congestion effects. A simple variant of the $LRTA^*$[6] algorithm for route computation is used to model a person's panic behavior.