

## MULTI-AGENT MODELING AND ANALYSIS OF THE BRAZILIAN FOOD POISONING SCENARIO

V. MYSORE,\* O. GILL, R.S. DARUWALA, M. ANTONIOTTI, and  
 B. MISHRA, Courant Institute, New York, NY  
 V. SARASWAT, IBM T.J. Watson Research Center, Yorktown Heights, NY

### ABSTRACT

The multi-agent modeling and analysis of catastrophic events raise many challenging problems since they involve a large, interacting, mobile population with complex behaviors. This research aims to address these problems through the analysis of simulations and to aid planning efforts for future catastrophic events through parameterized stochastic models covering the health care providers, emergency responders, and affected population. As a test case, we examine the massive outbreak of *Staphylococcus aureus* food poisoning that occurred in Minas Gerais, Brazil, in 1998 to demonstrate and evaluate our tools and techniques. In this incident, 8,000 people consumed contaminated food at a priest's ordination. Of these, 81 were admitted to intensive care units of 26 local hospitals after a triage, and 16 of them eventually expired. We capture the dynamics of such an outbreak by using two kinds of abstract agents — *hospital* and *person*, further augmented with *information* and *communication channels*. Hospital locations and current capacities are broadcast by the hospital to its patients and to persons with a radio and subsequently exchanged between neighboring persons. This “outbreak” model has been implemented in the Java version of Repast 3.0. Most attributes are scaled to be in the range of 0 to 1, with most behavior being probabilistic. We document the relative performance of the different simulations by using a range of parameter values for communication channels, personalities, and triage policies, to understand their combined effect on the overall survival rates. We also introduce the XSSYS trace analysis and model checking tool for answering complex temporal logic queries over Repast traces. We discuss how such simulation-based analysis can become a rigorous tool in aiding public health policy planning.

**Keywords:** Social simulation, catastrophe preparedness, emergency response, Repast

### INTRODUCTION

The computer modeling and simulation of catastrophic scenarios, when enhanced with sophisticated automated reasoning, promise to be a very valuable tool for developing public health policies and disaster management strategies. In the horrific wake of Hurricane Katrina that ravaged the State of Louisiana, it became doubly shocking as word spread very rapidly about the computer models that had accurately predicted many of the ramifications of such a disaster. Indeed, the Center for the Study of Public Health Impacts of Hurricanes of Louisiana State University had conducted extensive research on this topic and constructed elaborate models of such a scenario (see Heerden and Binselam 2004). While it is much less likely that other

---

\* Corresponding author address: Venkatesh P. Mysore, Courant Institute, 715 Broadway #1006, New York, NY 10012; e-mail: mysore@cs.nyu.edu.

simulation efforts can achieve such predictive fidelity, most catastrophe simulation projects (e.g., SEAS project of Chaturvedi et al. 2003, Project RESCUE of Mehrotra et al. 2004, and VISTA tool of Louie and Carley 2004) still focus on one of two nonoverlapping goals: *disaster prediction* and *disaster management*. In this paper, we do not even broach disaster prediction; instead, we focus on the analysis of simulations to aid planning efforts for future catastrophic events. We are part of the Large Scale Emergency Response (LaSER) research group of the New York University (NYU) Center for Catastrophe Preparedness and Response (CCPR), which is a partnership with the U.S. Department of Homeland Security and its Office for Domestic Preparedness. Catastrophe preparedness involves stocking and distributing resources to minimize fatalities, planning an emergency response strategy, and educating the general population. These desiderata will dictate, among other things, the distribution and use of available resources, and the means and nature of the information and instructions provided to the health care providers, emergency responders, and affected population (see Lasker 2004). This paper deals with these issues through the multi-agent modeling of catastrophic events that involve a large, interacting, mobile population with complex behaviors and goals.

We use the massive outbreak of *Staphylococcus aureus* food poisoning that occurred in Minas Gerais, Brazil, in 1998 (Do Carmo et al. 2004) to demonstrate and evaluate our tools and techniques. Although the fraction of fatalities (16/8,000) may not be regarded to be of catastrophic magnitude, the scenario is ideal for observing the effects of different instructions and policies on the behavior of the large affected population and the medical facilities. We capture the dynamics of such an outbreak by using two kinds of abstract agents — *hospital* and *person* — enhanced with *information* and *communication channels*. After exploring a number of simulation systems, this “outbreak” model has been implemented in the Java version of Repast (Collier et al. 2005). Most attributes are in the range of 0 to 1, with most of the behavior governed by random-number-based probabilities. We document the relative performance of the different simulations by using a range of parameter values for communication channels, personalities, and triage policies, to understand their combined effect on the overall survival rates. We also introduce the XSSYS trace analysis and model checking tool (Antoniotti et al. 2003) developed in our laboratory and show how it can answer complex temporal logic queries over Repast traces. We conclude by suggesting how such a schema provides a reasonable way of modeling, simulating, and analyzing other catastrophic scenarios as well.

## BRAZILIAN OUTBREAK

In 1998, a massive outbreak of *Staphylococcus aureus* food poisoning occurred in the rural town of Minas Gerais, Brazil, where around 8,000 individuals attended a Catholic priest’s ordination. The trace-back investigation implicated food preparers, who were culture positive for enterotoxigenic *Staphylococcus aureus*, as the source of contamination. However, it was the improper storage temperature of the food, which was prepared 2 days in advance, in the summer weather that allowed the optimal growth of bacteria and production of *Staphylococcus enterotoxin* (SE). Symptoms like intense nausea, emesis, diarrhea, abdominal pain, prostration, and dizziness were pronounced in less than 4 hours after consumption of the contaminated food in about half the population (~4,000). Almost half of them (~2,000) decided to proceed to one of the 26 nearby hospitals without letting the situation exacerbate further. However, this overwhelmed their emergency departments, forcing a triage. A triage, in medical parlance, refers to a set of policies to partition the vast number of patients into different groups (e.g., those requiring immediate intensive care, those requiring general hospitalization, and those requiring

only medication or saline). This process helps the hospital distribute the available resources optimally under the time constraints imposed by the prognosis of the disease. In Minas Gerais, 396 (~20%) people required admission after triage, and of these, 81 (~20%) required admission to the intensive care unit (ICU). Patients with improving health were discharged from the ICU within 7–10 days. A total of 16 (~20%) patients subsequently developed irreversible multi-system shock and expired while hospitalized. While people of all ages (1–86) attended the ordination, the 16 fatalities occurred only in the oldest (65 and above) and the youngest (5 and under) groups. The sex of the individual was found to have no influence on the clinical outcome among those treated in the ICU.

## MULTI-AGENT OUTBREAK MODEL

We capture the dynamics of such an outbreak by using two kinds of abstract agents: *hospital* and *person*. A hospital is an abstraction of any medical facility accessible in the area (26 in the Brazilian case), while a person is an abstraction of any individual who consumed the contaminated food (8,000 in the Brazilian case). The effect of the general population who did not attend the ordination is not modeled in our simulation. The model is then enhanced with *information* and *communication channels*, with the two vital pieces of information being the locations of the hospitals and their current capacities.

### Food Poisoning

The food poisoning is modeled by functions that describe the time variation of the person’s health, with and without treatment. Effectively, any “disease” can be modeled in terms of the (possibly time-varying) amount by which the affected agent’s “health” can deteriorate or recover with and without treatment, at each time-step of the simulation. The individual’s resistance or susceptibility to the specific disease is captured by a personalized variable, which modifies the disease-health-treatment functions. This can be used to abstract factors such as age, sex, health condition before food consumption, and genetic makeup. Probabilities are introduced to capture unpredictability and variability in real situations. We can use this simple but effective abstraction to model other conditions, such as Sarin gas attacks, radiation exposure, etc. Since the initial amount consumed and the dose/response relationship in human oral exposure to SE are unknown, the initial health of each person is assumed to be a random value in a meaningful range.

### People’s Behavior

The *persons* move toward their place of work from the site of food poisoning. Depending on their deteriorating health level and personality parameters, they choose to go to the one *hospital* they are initially aware of. Additional information is acquired by talking to neighboring agents. A time stamp of the information is maintained, so the persons update their knowledge only if more current information is available. Further, some persons are equipped with radios, which give them access to the current information about all the hospitals. People recompute the destination hospital toward which they should be moving on the basis on the distance to and the believed current capacity of each medical facility they are aware of. In addition, they always move toward the nearest free hospital, unless they are very sick and opt to go to the nearest

hospital, even if it is full. The complexity of the model is increased further with *personality* parameters, which capture whether an agent chooses to go to a hospital, talk to neighbors, accept the new information, or recompute the best hospital. *Group behavior* is captured by letting adjacent people moving toward the same destination wander less.

## Hospital Behavior

The *hospital* aims to admit every *person* who reaches its premises and invests its resources in the order of their admittance and proportional to their ill health. Hospital resources, consisting of infrastructure, beds, nurses, and doctors, are recovered when a patient is discharged or deceased; medical supplies, like drugs and saline, are irrecoverable. The hospitals also perform a local broadcast of complete current information to all persons who are admitted or waiting at their facility. The hospital model is enriched by identifying three different modes of operation — full, critical, and available — corresponding to the current amount of resources. With the *triage* policy in place, the hospital agent handles admitted persons as before. However, it admits new persons only if it has resources to spare (available mode). If it is operating in the critical mode, it admits only critically ill persons. No new persons are admitted in the full mode. With the *transfer* policy in place, admitted patients who have recovered reasonably are discharged earlier than usual and instructed to go to a different hospital if symptoms recur. In their place, critically ill persons who are waiting are admitted. Probabilistic parameters are used to capture the policies that govern the hospital's decisions on when to admit a new patient, in which order to treat the admitted patients, when to transfer a recovering patient to a nearby hospital, and which critically ill patient to admit in the vacancy created.

## ANALYZING THE OUTBREAK

Since the modeled system involves a large number of agents, uses a vast number of parameters, and attempts to capture the stochastic nature of the infection and behavior, traditional symbolic or algebraic analyses are not immediately possible. Instead, the analyst must resort to simulation-based analysis to obtain average performance statistics over a large number of trials. Combined with individual inspection of a small number of characteristic traces, evaluation of the relative merits of different emergency response strategies becomes possible. We use the statistics-based analysis tools provided by Repast and introduce the temporal logic trace analysis tool XSSYS.

## Numerical Results

Since the most significant aspect of the model is its extreme sensitivity and unpredictability, general average/comparative trends (as opposed to absolute values) in the death rate can be used to observe the effect of variations in parameters of interest (with the other dimensions fixed at justifiable values). We obtain trends (typically averaged over three runs) around the Brazilian scenario with 8,000 people and 26 hospitals leading to a death rate of 0.2%.

### *Effect of Hospital Resources, Communication, and Grid-Size on Death Rate*


We first observe how the number of deaths varies with hospital resources (Figure 1). Shown there are the plots for a  $250 \times 250$  grid and an  $800 \times 800$  grid, with communication enabled and disabled, and with no triage policy implemented. From the plot, we observe that the number of deaths clearly declines when hospitals have more resources, since each hospital is able to allocate more resources (treatment) per person. Also note that in a small grid ( $250 \times 250$ ), where hospitals have few resources, communication works against our model. This is because people converge to the nearest hospitals, exhausting their resources quickly. By the time the hospital runs out of resources and turns people away, they are too sick to survive a trip to the next hospital. However, when the hospitals have plenty of resources, the difference in survival rates is negligible when communication is used versus when it is not. In the  $800 \times 800$  cases, the difference in distances between the closer and farther hospitals is much greater. Hence, it works to a person's benefit to communicate and obtain information about nearby hospitals.

### *Effect of Number of Hospitals, Triage, and Grid Size on Death Rate*


Next we analyze how the number of hospitals affects the number of deaths by using plots (Figure 2) for a  $300 \times 300$  grid and an  $800 \times 800$  grid, with the triage policy enabled and disabled. We first note the expected phenomenon: increasing the number of hospitals decreases the death rate, since there are fewer patients per hospital. We also note a slightly higher death rate when the grid size is larger because the average distance to a hospital is longer: people reach the hospitals when they are sicker, and more persons are not able to survive the journey. More important, this figure leads to a dramatic conclusion: the triage policy, as interpreted in the model, always works against the people. The failure of the triage policy can be attributed to a key aspect of the food-poisoning health function: a healthy person is just as likely to worsen as an already unhealthy person. Thus the patients who were discharged slightly early because of the critically ill people who were waiting end up falling sick again, and the critical ill persons themselves seldom recover. Second, the health of people who are refused admission (because they are not critically ill or because the hospital is full) worsens during their trip to a different hospital. The net effect is that the hospitals have to treat sicker people. This suggests that it is wiser for people to reach the nearest hospital, and then for the hospitals to have a system of redistributing their resources (i.e., moving equipment and doctors, as opposed to moving patients).

### *Effect of Number of People, Grid Size, and Initial Pattern on Survival Rate*


Next, we observe how the number of people affects the fraction of people who survive by using plots (Figure 3) for a  $300 \times 300$  grid and an  $800 \times 800$  grid, with communication enabled and the triage policy disabled. We also inspect the effect of people starting at random positions in the grid as opposed to being concentrated at a location. From these plots, we again observe the expected trend: as the number of people increases, the fraction of people who survive declines. Similarly, the  $800 \times 800$  grid results in a slightly larger percentage of the people dying because the average distance to the nearest hospital is longer. The difference in survival percentages for the concentrated and the random initial positions is not statistically significant. This can be understood as the average person's starting point's distance to the nearest hospital being roughly



**FIGURE 1** Effect of hospital resources, communication, and grid size on death rate



**FIGURE 2** Effect of number of hospitals, triage, and grid size on death rate



**FIGURE 3** Effect of number of people, grid size, and initial pattern on survival rate


the same in both cases. However, the number of initial neighbors in the distributed case must be sufficient to supply the required information about the nearest hospitals.

### Trace Analysis in XSSYS

The XSSYS temporal logic trace analysis system can answer linear temporal logic (LTL) queries about the time course behavior of a set of traces. It was developed originally as a part of Simpathica for simulating and analyzing biochemical pathways. XSSYS allows the user to formulate queries about multiple traces in temporal logic or English (via a natural language interface). The person and hospital traces of Repast can be read by using XSSYS. These traces reveal very insightful aspects of the behavior of persons and hospitals and serve as a good starting point for coming up with new policies to be tested. Complex temporal queries linking different traces can help in discovering finer truths about the underlying dynamics of the system. In this section, we demonstrate the XSSYS trace analysis tool in some simple examples.

#### Time-trace of a Person

The variation of a person's health with time (in this case, *Person-78*) during the course of a simulation is plotted in Figure 4. XSSYS plots this curve by using data imported from Repast in the *btd* format by using the *PtPlot* tool. In addition to the health level (*HealthLevel*), the person's current location (*x*, *y*) and destination (*destx*, *desty*) are plotted. To indicate when the person actually received treatment, a Boolean value *admitted* is also plotted.



**FIGURE 4** Time-trace of a person

#### *Time-trace of a Hospital*

In the case of a hospital, we plot the depletion of resources (*HospitalResources*) with time (Figure 5). The number of people admitted and the number of people waiting indicate the stress on the hospital (in this case, *Hospital-1*). The successful creation of vacancies by early discharge and their filling by critically ill persons awaiting treatment are also presented.

#### *Temporal Logic Analysis*

Temporal properties of these traces can be analyzed by formulating queries in linear temporal logic by using the operators *Eventually* (sometime in the future) and *Always* (henceforth in the future). In the specific case being demonstrated (Figure 6), the traces of *Person-13* and *Person-113* are being compared. *Person-113* is seen to have a consistently better *HealthLevel* than *Person-13*, although both their *HealthLevels* are dropping. *Person-113* is also seen to have reached the destination hospital, while *Person-13* has not.




FIGURE 5 Time-trace of a hospital




FIGURE 6 Temporal logic analysis in XSSYS

## DISCUSSION

The Brazilian food-poisoning scenario proved to be a considerably complex problem, which had all the essential elements of a typical catastrophic scenario: a large number of agents (8,000 + 26), agents of different types (persons and hospitals), external factors governing the time evolution of the agents' features (effect of food poisoning on health), mobility (persons), mutual interaction (within persons, and between persons and hospitals), and multiple communication channels (talking, broadcast, and radios).

Repast proved to be sufficient to model and simulate the Brazilian food poisoning scenario. The analytical capabilities were enhanced by feeding its output to XSSYS. Despite the extreme parameter sensitivity of the model, we were able to explore the effectiveness of different emergency response strategies and catastrophe preparedness policies. The complexity and unpredictability of the model, because of the vast number of parameters, became apparent very quickly. Our model was able to capture the reported statistics to a reasonable extent, and it elucidated different conditions that could have led up to them. Factors that could have increased or decreased the number of fatalities also became evident. More specifically, the results showed that the distance the people need to travel to reach the hospital greatly determines how many people survive. We also observed that the survival rate increases when either the resources each hospital has or the number of hospitals increases, and that the survival rate decreases when the number of people increases. When the average distance to the nearest hospital is almost the same, there is almost no difference in survival rates between concentrated and random initial patterns. We found that communication among people about hospitals is beneficial when the difference in distances to hospitals is substantial, but it is harmful when all hospitals are close by and have few resources. We also found that our triage system harms the survival rate, since it is better to keep patients at a hospital, even if it has low resources, rather than have them transfer to another hospital and then having to treat a sicker person. The emergence of such interesting unanticipated behaviors already suggests a potential utility of such simulation-based analysis tools.

Many additional enhancements to the outbreak model to make it more realistic are possible. We might need to switch the environment to a real city. Transportation constraints and modes, roads, subways, and other geographical information might need to be incorporated. The moment these additional constraints emerge, we will need to model the agent's transportation choices. For example, Raney and Nagel (2004) describe a framework for running large-scale multi-agent simulations of travel behavior on the basis of each agent's "plan" of activities, times, and preferred modes of transport. However, as described by Sono and Ishibashi (2004), the change in the transportation choices after a disaster will need to be worked into the plan, with commuters and noncommuters having to be treated differently (a rather simple situation, which nonetheless seems to have had a major impact in the Katrina disaster). A somewhat complex model of this nature will endow each agent with a current-mobility variable, which decreases with a decrease in the agent's health, increases if the agent is being helped by a neighbor, and decreases if the agent is helping a neighbor.

We will need to add social networks at various levels (families, friends, etc.) and the social characteristics of subsets of the population to model the cultural differences in response behavior. A good example of the application of social judgment theory appears in the work on group attitude emergence via assimilation and contrast effects as described by Jager and Amblard (2004). The benefits of cooperation could be captured by increased mobility and

information, while moving in groups. We could also add social infrastructure, like first responders, volunteer-based relief organizations, and law enforcement officers. Also, some of the people who consumed the contaminated food could belong to these groups, thus complicating the interaction dynamics even further.

We could also add more detailed models of communication and information exchange. For instance, the logic-based framework for handling messages and belief-state changes discussed by Perrussel and Thevenin (2004) could be combined with ideas from the work on the geographical divergence of knowledge via interactive-learning-based diffusion by Morone and Taylor (2004). This could prove useful in capturing the realistic transmission and accumulation of information during calamities. We could incorporate into the model long-distance 1-to-1 and 1-to-many communication channels, where 1-to-1 channels are between persons via cell-phone and 1-to-many are from authorized broadcasters to equipped receivers. We could model the ability to give instructions and the ability to receive instructions separately. Similarly, there could be a difference in the transmission of different kinds of information (e.g., the location of the nearest hospital, measures to use to slow down the progression of the sickness, instruction to proceed to a hospital). (See the work of Lawson and Butts [2004] on the propagation of rumors and information in crisis contexts.)

The food poisoning in itself could have been modeled differently. For instance, the spread of *Mycoplasma pneumoniae* via interaction between patients and caregivers is modeled by using network theory by Meyers et al. (2003). Similarly, Rahmandad and Sterman (2004) analyze the pros and cons of agent-based modeling versus differential equation modeling for contagion modeling. Although the work of Eidelson and Lustick (2004), who developed a stochastic agent-based model, VIR-POX, to explore the viability of available containment measures as defenses against the spread of smallpox, is similar to the Brazilian scenario analysis, it is different in its approach and goals.

On the pure computational side, the biggest challenge is in scaling up to a very large-scale simulation through parallelization, abstraction, hierarchy, and other strategies. We are working on enhancements to XSSYS to improve its expressivity and power. We also need to investigate the applicability of other formal reasoning techniques, such as probabilistic reasoning (Xiang 2002) and probabilistic argumentation systems and causal analysis (see the WIZER tool of Yahja and Carley 2004). We could treat the estimation of the triage policy parameters (e.g., the health level at which a person who is waiting gets deemed as critically ill, or the health level at which a recovering patient may be discharged to create a vacancy) as an optimum-value computation problem. From a practical utility point of view, we need to identify a way of describing the simulations in a manner that is formal and accurate enough to create a meaningful simulation but simple enough for a nonprogrammer to use. We are in the process of compiling a survey of approaches to model and analyze catastrophic scenarios. Our goal is to first extend this modeling and analysis approach from the Brazilian food-poisoning example to other scenarios. For example, the effects of several people independently consuming botulinum-contaminated milk at their homes (following the scenario investigated by Wein and Liu [2005]) could be modeled by a different health-modulation curve, and with people starting at their homes as opposed to congregating at a church. We would additionally need to model the transmission of the instruction to not consume any more contaminated milk. Eventually, we hope to develop and demonstrate the tools and technologies necessary for such simulation-based analysis to provide a rigorous yet user-friendly approach for exploring assumptions about public health policies in catastrophe preparedness and emergency response.

## ACKNOWLEDGMENT

The research reported on in this paper was funded by a LaserMed grant from NYU's CCPR.

## REFERENCES

Ancel Meyers, L., M.E. Newman, M. Martin, and S. Schrag, 2003, "Applying Network Theory to Epidemics: Control Measures for *Mycoplasma pneumoniae* Outbreaks," *Emerging Infectious Diseases* 9(2):204–210, Feb.

Antoniotti, M., F. Park, A. Policriti, N. Ugel, and B. Mishra, 2003, "Foundations of a Query and Simulation System for the Modeling of Biochemical and Biological Processes," in *The Pacific Symposium on Biocomputing: PSB 2003*, (Editors R.B. Altman, A.K. Dunker, L. Hunter, T.A. Jung, and T.E. Klein), pp. 116–127, World Scientific, Jan.

Chaturvedi, A., S. Mehta, and P.L. Drnevich, 2003, "Measured Response 2003," invited white paper for the Purdue e-Business Research Center (PERC) and Homeland Security Institute meeting.

Collier, N., T. Howe, R. Najlis, M. North, and J.R. Vos, 2005, "Repast 3.0: Recursive Porous Agent Simulation Toolkit;" available at <http://repast.sourceforge.net/>.

Do Carmo, L.S., C. Cummings, V.R. Linardi, R.S. Dias, J.M. De Souza, M.J. De Sena, D.A. Dos Santos, J.W. Shupp, R.K. Pereira, and M. Jett, 2004, "A Case Study of a Massive Staphylococcal Food Poisoning Incident," *Foodborne Pathogens and Disease* 1(4):241–246, winter.

Eidelson, B.M., and I. Lustick, 2004, VIR-POX: An Agent-based Analysis of Smallpox Preparedness and Response Policy," *JASSS* 7(3), June 30.

Jager, W., and F. Amblard, 2004, "A Dynamical Perspective on Attitude Change," NAACSOS Conference, Pittsburgh, PA, June.

Lasker, R.D., 2004, "Redefining Readiness: Terrorism Planning through the Eyes of the Public," The New York Academy of Medicine, New York, NY.

Lawson, G.C., and C.T. Butts, 2004, "Information Transmission through Human Informants: Simulation," *CASOS'04*.

Louie, M.A., and K.M. Carley, 2004, "VISeualization of Threats and Attacks (VISTA): A Decision Support Tool for Urban Threat Environments," *CASOS'04*.

Mehrotra, S., C. Butts, D. Kalashnikov, N. Venkatasubramanian, R. Rao, G. Chockalingam, R. Eguchi, B. Adams, and C. Huyck, 2004, "Project RESCUE: Challenges in Responding to the Unexpected," *SPIE*, Vol. 5304.

Morone, P., and R. Taylor, 2004, "Small World Dynamics and the Process of Knowledge Diffusion: The Case of the Metropolitan Area of Greater Santiago De Chile," *Journal of Artificial Societies and Social Simulation*, Vol. 8.

Perrussel, L., and J.M. Thevenin, 2004, "A Logical Approach for Describing (Dis)belief Change and Message Processing," in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems*.

Rahmandad, H., and J. Sterman, 2004, "Heterogeneity and Network Structure in the Dynamics of Diffusion: Comparing Agent-based and Differential Equation Models," *CASOS'04*.

Raney, B., and K. Nagel, 2004, "An Improved Framework for Large-scale Multi-agent Simulations of Travel Behavior," in *Proceedings of Swiss Transport Research Conference (STRC)*, Monte Verita, Switzerland; available at [www.strc.ch](http://www.strc.ch).

Sono, K., and K. Ishibashi, 2004, "Constructing Commuters' and Non-Commuters' Action Choice Model after a Great Earthquake," *CASOS'04*.

van Heerden, I.L., and A. Binselam, 2004, "Hurricane Pam Exercise — July 2004, Adcirc Storm Surge Animations and Related Information," Center for the Study of Public Health Impacts of Hurricanes; available at [http://hurricane.lsu.edu/floodprediction/PAM\\_Exercise04/Introduction/Introduction.htm](http://hurricane.lsu.edu/floodprediction/PAM_Exercise04/Introduction/Introduction.htm).

Wein, L.M., and Y. Liu, 2005, "Analyzing a Bioterror Attack on the Food Supply: The Case of Botulinum Toxin in Milk," *PNAS* 102(28):9984–9989, July 12.

Xiang, Y., 2002, *Probabilistic Reasoning in Multiagent Systems: A Graphical Model's Approach*, Cambridge University Press.

Yahja, A., and K.M. Carley, 2004, "The Design of Automated Validation and Explanation for Large-scale Social Agent Systems," *CASOS'04*.

