Proceedings of the 2007 Winter Simulation Conference
S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, eds.

PANEL: AGENT-BASED MODELING OF MASS EGRESS AND EVACUATIONS

Douglas A. Samuelson

Serco
2650 Park Tower Drive, Suite 800
Vienna, VA 22180

Matt Parker

ANSER
2900 S. Quincy Street, Suite 800
Arlington, VA 22206, U.S.A.

Austin Zimmerman

Homeland Security Institute
2900 S. Quincy Street, Suite 800
Arlington, VA 22206, U.S.A..

Stephen Guerin

Redfish Group
624 Agua Fria Street
Santa Fe, NM 87501, U.S.A.

Joshua Thorp

Redfish Group
624 Agua Fria Street
Santa Fe, NM 87501, U.S.A.

Owen Densmore

Redfish Group
624 Agua Fria Street
Santa Fe, NM 87501, U.S.A.

Pat McCormick

Alpha Informatics Ltd.
8809 Sudley Road, Suite 201
Manassas, VA 20110, U.S.A..

Tom McCormick

Alpha Informatics Ltd.
8809 Sudley Road, Suite 201
Manassas, VA 20110, U.S.A..

ABSTRACT

We will discuss several recent advances in agent-based modeling, with applications to mass egress and wide-area evacuations following disasters. These advances include: efficiency improvements in specifying people's identification and selection of exit routes, making it possible to handle up to 70,000 people-agents; incorporation of new crowd movement depiction ("Continuum Crowds") for greatly increased realism; efficient real-time visualization depiction of the resulting behaviors; and including the effects of communication and direction, on scales ranging from individual facilities to metropolitan areas.

1 INTRODUCTION

For the Science and Technology Directorate of the Department of Homeland Security, a team from the Homeland Security Institute, along with Redfish Group, developed agent-based simulation models of mass egress from a stadium and a subway station following one or more detonations of improvised explosive devices (IEDs). These models focus on representations that support evaluation of different proposed ways to facilitate evacuation and reduce casualties.

The stadium model represents PNC Park, the major league baseball stadium in Pittsburgh, Pennsylvania. Anti-IED countermeasures modeled included improved guidance to exits, egress onto the playing field, and shock- and shrapnel-absorbing baffles inside the spiral stairways that provide the largest-volume routes of egress. This model depicts up to 70,000 people, at least an order of magnitude more than in previous real-time agent-based simulation models.

The subway model represents one level of the Metro Center subway station in Washington, D.C. This model incorporates new computational portrayals of explosions and crowd movement. To the extent that experiments are possible, these approaches appear to conform more realistically to real events than the methods used previously in agent-based models of such phenomena. Countermeasures represented included improved guidance to exits and baffles on the platform to absorb shock and shrapnel.

To support an analysis of DoD response to a nuclear terrorist attack on Washington, DC, a study team from

Alpha Informatics, Limited (AIL) and SPARTA, Inc. examined the impact of alternative communications architectures on the outcomes of the scenario.

2 THE MASS EGRESS PROBLEM

If one or more IEDs were detonated in a large public facility in the U.S., the attempt by hundreds or even thousands of people to leave the facility quickly, all at the same time, could result in many casualties secondary to those caused by the explosions. If the explosions or some additional action produced a toxic debris plume, as well, this would produce additional hazards to the people trying to leave. It is worthwhile to understand what countermeasures have the greatest potential to reduce secondary harm.

2.1 Stadium

The study team selected PNC Park because the team determined that the necessary cooperation and data could be obtained relatively easily and because the City of Pittsburgh had recently conducted a live exercise to assess potential response to a biochemical incident.

Although Redfish Group began building the model in NetLogo, the team soon realized that real-time execution would not be possible on this scale. Accordingly, the team switched to Processing, an add-on to Java, for this model.

The model displays a depiction of the stadium in "dots-and-wires" form: that is, the people appear as colored dots within a somewhat stylized but still accurate picture of the stadium. The user can pre-specify explosions before the run or introduce them interactively by using the mouse to position the cursor to a location and then pressing the "b" key (for "bomb"). Explosions are depicted as spherical events, with specified probabilities of death or injury out to certain radii. After the first explosion, people are depicted as moving toward exits, away from the explosion site, at various rates. If a subsequent explosion occurs, blocking intended paths of egress, people are represented as altering their movement to seek a different exit.

People-agents employ a least-cost-path assessment to decide which exit, of those they know about, to attempt to use. A conventional way to provide distance or cost information to agents is a "flood fill" algorithm, in which the depth of simulated water at each location, following a large in-flow through a given exit, is inversely proportional to the distance to that exit from that location. Obstacles can then by portrayed as increased distances through the locations they occupy, and people-agents adjust their paths accordingly. The problem is that doing these calculations for the whole venue, for each person-agent, at each time step is prohibitively compute-intensive. By having each person-agent identify a few ribbons of locations on the way to each exit, however, and by then updating only the locations in a person-agent's identified paths, the team found that real-time performance for a large number of people-agents could be obtained.

The output includes a real-time, movie-like (albeit somewhat abstract) depiction of behavior, some summary statistics (on screen, written to a file, or both, at the user's option), and an output file. This output file can be submitted, in turn, to a rendering program that will produce a QuickTime movie, with a much more realistic visual portrayal of the stadium and the people. The problem is that such renderings are impossible to produce on a single computer in less than several days: that is, the rendering into a movie is possible, but unacceptably slow for most purposes. Therefore, the user can either obtain the rendering software and produce movies of selected runs, allowing the time required to do so, or submit the output file either to a "rendering farm" provider or to HSI and Redfish, who will in turn submit the file to a rendering farm to produce the movie. A rendering farm translates files into a form that can be processed by massively parallel arrays of large numbers of processors, transmits appropriate components of the data to those processors, and assembles the results. This rendering takes less than two hours, rather than several days.

A significant advantage of this approach is that a decision-maker with very little knowledge of or interest in simulation methods can observe the simulated patterns of crowd movement under various sets of assumptions and see differences among the patterns. Thus a facility manager can also assess quickly whether proposed changes in structure or in evacuation procedures are likely to help much.

2.2 Subway

The subway model is in NetLogo, as the maximum number of people in the subway station — around 2,000 — is, at least for a modest level of complexity of behavior, within the scale NetLogo can handle with acceptable speed of execution. At this time, the team is still experimenting with different depictions of crowd behavior, as the movements and interactions in this scenario are more complex than in the stadium due to the limited number of exits and the implications of alternative choices. For example, if people try to escape the station by jumping into the track bed and fleeing along the tunnel, results depend on how quickly the transit authority can shut down power to the third rail. The team is also finding, however, that adding too many choices of crowd behavior could push the model beyond what NetLogo can handle with acceptable execution times. The determination of appropriate model scope continues, therefore, with development and experimentation.

People are represented as agents. Locations are represented as patches, and each person has a grid of patches around him that he perceives. Agents assign a potential to each patch in their grid, based on how crowded that patch looks and how undesirable it appears for other reasons (such as, a bomb explosion has taken place there, or there are dead bodies there.) Agents try to progress toward known exits without entering excessively high-potential patches.

In addition, the team learned of and employed a new algorithm, "Continuum Crowds," to provide more realistic representation of interactions among people-agents. In this algorithm, people-agents observe a continuous, smoothed assessment of how crowded the nearby patches are, so they anticipate crowding some distance ahead of them. The developers of the algorithm claim that it accurately reflects real crowd behaviors that other algorithms do not capture.

Bomb explosions are represented as a wave of shock agents, propagated in all directions, and swarms of shrapnel agents flying level, at velocity diminishing over distance. If a shock agent hits a person, the person notes the force with which he was hit. Shock agents reflect off hard surfaces and may hit the same person(s) more than once. If a shrapnel agent hits a person, it stops and the person notes that he is injured. Agents use a probability table to decide whether the combined effects of shrapnel and shock cause them to die. Dead people-agents remain on the patch where they died.

When too many agents try to enter the same patch, a probability table determines which ones get trampled. Trampled people remain, injured or dead, on the patch where they were trampled.

Here also, as in the stadium, technologies of particular interest are barriers that can absorb shock and shrapnel, and improved methods of directing evacuation. In the subway, unlike the open-air stadium, overpressure from the explosion is a serious effect, because the location is enclosed by stiff, reflective materials such as steel-reinforced concrete. The HSI team investigated properties of such materials; this aspect of the modeling necessarily involved some engineering judgment and conjecture. In the end, the team decided that modeling shock waves as reflecting, with diminishing velocity, from hard surfaces would suffice to approximate the combined effects of overpressure and reflected concussion.

2.3 Preliminary Findings

Runs of the two models to date yielded some interesting findings, including some that were counter-intuitive:

- Better guided egress definitely accelerates evacuation but — contrary to our initial expectations — does not appear to decrease casualties from trampling. In fact, expedited egress may *increase* trampling. For the stadium, this effect was greater when people were not permitted to run onto the field. While this may be an artifact of the programming, the team deems it more likely that it is a real effect, as better information about which way to exit induces more crowding near the good exits.
- Letting people run onto the field was the most beneficial measure tested in the stadium. (Note that different conditions, such as a chemical or biological plume moving toward the field, or the need to land med-evac helicopters there, would change this conclusion.)
- In the subway, knowing whether there is fire is critical: which way people should flee depends heavily on the nature of the event.
- As the team expected, baffles cut casualties from an explosion near them, typically by as much as 40 percent. However, given that the baffle is in place before the bomber decides where to detonate the explosives, there might be *no* net benefit from baffles in actual operation, because the bomber can simply detonate somewhere else.
- The results are critically dependent on assumptions. In some early runs on the stadium, a relatively minor change in crowd movement logic (an update to the all-or-nothing trampling logic and to how dense the crowd had to be for trampling to occur) reversed the preliminary finding about which egress protocol (guidance to exits versus confusion) produced more casualties.
- Not all the dependencies on assumptions are obvious. It took the team some time and considerable thought to distinguish between plausible findings and probable programming artifacts in a few cases.
- Because of the importance and subtlety of assumptions, the models are good enough to point out important issues, but too assumption-dependent to be good for real-time incident management.
- Standard validation methods have limited applicability to these models. Statistical significance tests are based on the assumption that only random variation — essentially sampling error — and the effect being modeled contribute to total variation. In this analysis, there are many more or less arbitrary assumptions about behavior and materials, and changing those assumptions would most likely cause substantial changes in the results. Hence uncritical application of standard estimation procedures would yield spurious precision.

- The models are highly sensitive to factors about which real-life data are scarce, and the available data are insufficient to support precise estimates of effects. In particular, there are few if any sets of high-quality real-event data about how the nature and urgency of a threat affects people's behavior. Therefore, the models are useful for raising questions and suggesting likely advantages of some mitigating measures, both in assessing technologies and in incident planning, but not for real-time prediction and incident management.

3 MASS EVACUATION

The Alpha Informatics (AIL) – SPARTA team analyzed the likely effects of a surface nuclear detonation, with a 10-kiloton TNT-equivalent yield, near the White House. To do this, the team used the Communications Architecture Support Tool (CAST), a discrete event simulation developed by AIL, and the Systems Effectiveness Analysis Simulation (SEAS), an Air Force toolkit agent-based complex adaptive systems simulation developed and maintained by SPARTA. The agent-based component of the analysis provides a depiction of the ways in which people will probably try to leave the area, the likely effect of obstacles and (known) contamination, and the efficacy of different checkpoint / decontamination / containment protocols and facilities.

After discussions with several civil First Response organizations, the team developed a detailed representation of both the civil First Responders and DoD elements included in the scenario, and implemented this representation in SEAS. In addition, the team developed a simulation incorporating the civil First Responder 800 MHz radio system, the civil PSTN, and the DoD communications capabilities in CAST to examine the flow of mission critical information in the scenario. The concepts driving both of these efforts were based on, and synchronized with, the response plan prepared by the Washington Metropolitan Council of Governments, the National Incident Management System, applicable local support agreements, and DoD organizational capabilities.

The team also had to develop a concept of operations (CONOPS) to drive the simulation. The team's CONOPS included the establishment of four regional support complexes, one at Fort Meade in Howard/Anne Arundel Counties, Maryland, one in Montgomery County, Maryland, one at Andrews AFB in Prince Georges County, Maryland, and one at Fort Belvoir, in Fairfax County, Virginia. These locations correspond well with the jurisdictional breakout of the organizations which would be collaborating on the response, and can be managed by a chain of command which fallows jurisdictional lines. Each of the areas identified has large open areas where temporary shelter can be set up, medical facilities where treatment and support can be provide, access to the road network for ground transport, and an airfield capable of operating C-130 class aircraft to provide air resupply and evacuation. In addition to these facilities, traffic control, and immediate decontamination and triage points would be established to receive, screen and route evacuees as needed. Finally, radiological reconnaissance and elements and field command and control facilities would be established to manage the operation, under the coordination of a centralized operations center.

The tools used to perform the simulation were the Communications Architecture Support Tool, which models the flow of information through a target communications architecture, and the Systems Effectiveness Analysis Simulation, or SEAS, an agent based, complex adaptive systems simulation which models the actions of the entities described in the scenario. SEAS is sensitive to the flow of information providing situational awareness to the entities, and is therefore particularly suited to an assessment of how the varying communications architectures would allow the DoD response to work with the area First Response elements.

The simulation posed a number of challenges:

- The problem is large. There are 180 to 225 individual organizations to be modeled, plus about 500,000 permanent residents and about 470,000 commuters.
- The simulation must account for the behavior of civilians under disaster conditions.
- Civilian management has primacy at the state and local level, so there would be ad hoc chains of command, with DoD and non-DoD elements intertwined.
- Communications architectures are generally *not* easily interoperable. They are, however, well documented and well understood.

The resulting simulation provided insights into the impact of communications interoperability on the outcome of the scenario, as well as operational planning insights which may be of use to first response planners. It revealed that the issues of interoperability are likely to be significant, but that, with advance planning, there are already capabilities extant which can facilitate the required interoperability. It also indicated that, for the period of the first 36-48 hours after an event such as that postulated in the scenario, Metro area First Responders have the necessary training, capabilities and are present in sufficient number to begin effective response and consequence mitigation.

4 GENERAL COMMENTS

These modeling efforts and the accompanying analyses extended the range of situations to which agent-based models have been shown to be applicable, and raised many interesting questions about appropriate responses. They illustrate the difficulty of the problems that can now be tackled, and some of the ways to work around the challenges we encountered. We believe these studies clearly show the potential for agent-based models to assist in planning and preparedness exercises, and thereby to support policy and architectural decisions to mitigate the expected effects of disasters.

ACKNOWLEDGEMENTS

The HSI-ANSER-Redfish team gratefully acknowledges the contributions of Loren Miller, Brett Steele, H. Ric Blacksten, Stacey Pommerenck, and Lonnie Johnson, and the helpful comments and suggestions of many colleagues, especially Peter Hahn, J. McGregor Smith, Rob Axtell, Rick Kohout, and Sarah Maloney.

AUTHOR BIOGRAPHIES

DOUGLAS A. SAMUELSON is a Principal Decision Scientist for Serco, in Vienna, Virginia. He was a Senior Analyst for Homeland Security Institute (HSI) from December 2005 through June 2007 and directed the facilities mass egress project described here. He specializes in probability modeling and simulation, emphasizing public policy applications. He has a B.A. in statistics from the University of California, Berkeley, and M.S. and D.Sc. in operations research from George Washington University.

MATT PARKER is an analyst for ANSER, leading and supporting computer projects for ANSER and HSI. He has a B.S. degree in mathematics from Purdue and a M.S. in computational science from George Mason University.

AUSTIN ZIMMERMAN joined HSI as an associate analyst in June 2006, following her graduation from M.I.T. with dual B.S. degrees in physics and management science.

STEPHEN GUERIN is President of Redfish Group, a R&D consultancy based in Santa Fe, New Mexico applying the emerging science of *Complex Adaptive Systems*. His work centers on visualization, modeling and the design of self-organizing systems. He also lectures on agent-based modeling and visualization as a faculty member of Santa Fe Institute's Complex Systems Summer School.

JOSHUA THORP is a principal analyst with Redfish Group, with several years of experience in complexity modeling with a variety of applications. Lately he has concentrated on modeling traffic evacuation in a wild fire scenario effecting Santa Fe neighborhoods. He has a B.S. in Computer Science with a concentration in Mathematics from the Cornell University School of Engineering.

OWEN DENSMORE is a principal analyst with Redfish Group. He has extensive experience in the computer industry, including work for Xerox, Apple, and Sun Microsystems. He holds 6 patents in computing areas. His current research is in Supply Chain Management, Multi-Agent Simulation, and Modern Optimization Techniques. He earned his B.S. in Math and Physics from Georgia Tech in 1964, spent two years in Africa with the Peace Corps, and earned his M.S. in Physics from Syracuse University in 1968.

PAT MCCORMICK is CEO of Alpha Informatics, Limited. She holds a B.S. in Computer Science from the US Military Academy, West Point, and an M.S. (with Distinction) in Operations Research from George Mason University. Her Web page is <www.inforead.com>.

TOM MCCORMICK is a Senior Analyst for Alpha Informatics, Ltd. He holds a B.A. in History from Dickinson College, an M.S. in Systems Management from Troy State University, and an M.A. in International Relations / Strategic Studies from Boston University. He is also a graduate of the US Army Command and General Staff College and the Joint Forces Staff College.