Agent-Based Simulations of Mass Egress After Improvised Explosive Device Attacks

Douglas A. Samuelson
Serco
samuelsondoug@yahoo.com

Matthew Parker
ANSER
matthew.parker@anser.org

Austin Zimmerman
Homeland Security Institute
(formerly)

Loren Miller
Homeland Security Institute
(formerly)

Stephen Guerin
Redfish Group
stephen.guerin@redfish.com

Joshua Thorp
Redfish Group
jthorp@redfish.com

Owen Densmore
Redfish Group
owen.densmore@redfish.com

ABSTRACT

For the Science and Technology Directorate (S&T) of the U. S. Department of Homeland Security, we developed agent-based computer simulation models of mass egress from a stadium and a subway station following one or more attacks with Improvised Explosive Devices (IEDs.) Anti-IED countermeasures we modeled included improved guidance to exits, baffles to absorb shock and shrapnel, and, for the stadium, egress onto the playing field. We found improved real-time information systems that provide better guidance to exits would substantially expedite egress and could reduce secondary (trampling and crush) casualties. Our results indicate that models like these can be useful aids to selecting countermeasures, and for training, preparation and exercises. We also discuss the unusual problems such models pose for real-time event management and for validation and evaluation.

Keywords

Agent-based simulation, mass egress, emergency preparedness, improvised explosive devices

INTRODUCTION

This analysis by the Homeland Security Institute (HSI) was part of a larger project intended to help S&T prioritize research efforts to address national needs for improved countermeasures against Improvised Explosive Device (IED) attacks within the U.S. To prepare for the modeling task, the HSI team first identified metrics, primarily number of casualties and speed of egress. Data on cost of response and secondary damage to the public were not available with sufficient accuracy and detail to be useable. Other metrics may also be appropriate, but HSI deemed these sufficient for this test of whether the model-based evaluation approach is useful. There followed a thorough discussion, with S&T and numerous other parties, of scenarios, possible countermeasures, bomb squad commanders' selection of unmet needs, and technology providers' information about new and proposed countermeasures

HSI next used model building tools, specifically agent-based simulation, to analyze security screening operations, concepts of operations, and the interrelationships among various input parameters. HSI originally contemplated using discrete-event simulation, also known as Monte Carlo simulation. This approach uses stylized, abstract representations and broad simplifying assumptions. One of the most important limitations of this approach is that individuals in the system must be treated as behaving more or less identically: for each class of individuals of a given type (there may be more than one type), the modeler specifies some probability distributions of how individuals will act at certain places in the system, and every individual in a class acts in accordance with the same probability distributions.

The scope and focus of the modeling changed over the course of the project, and the choice of method changed accordingly, from traditional discrete-event simulation to agent-based simulation. The team also considered analytical approaches but rejected them because of the difficulty of analytically modeling large numbers of transient events and the consequent difficulty of addressing "what-if" questions without substantial additional data collection and modeling.

HSI used software that is readily available commercially; in fact, it is open source, so the modeling executable as well as the source code for the modeling tools is readily and freely available.

HSI planned to rely on outside assessments, where available, of the capabilities of technologies in order to evaluate the relevance and efficacy of proposed technologies embedded in the threat scenarios. As in the model building and analysis, this activity changed in focus and scope as the project progressed, and the focus group meetings to select technologies did not produce clear results. HSI and the sponsor eventually chose a few representative technologies and evaluated their effects based on best scientific and engineering judgment.

After delivery of preliminary results, HSI continued to examine additional variations and alternatives. The models proved to be very sensitive to small changes in behavioral assumptions, rendering additional, more precise sensitivity analysis unnecessary.

Because egress is the more challenging problem, relative to the effects of screening in ingress, from the modeling standpoint, and because egress has not been the primary focus of previous studies, the HSI team concentrated much of its early analysis on technologies to direct and improve egress. In addition, the sponsor decided, after an early demonstration, that models such as these could themselves be among the more useful technologies for planning and event management. Therefore, the HSI team and the sponsor agreed to focus on technologies to improve egress and to mitigate effects of attacks directed at egress routes.

HSI engaged Redfish Group, a software and visualization company in Santa Fe, New Mexico, to construct the stadium model. While Redfish Group is among the world's leaders in agent-based models of crowd phenomena, this project exceeded, in both size and complexity, anything Redfish Group had previously done. The original plan to build this model in NetLogo, for ease of use and learning, gave way when the resulting model ran unacceptably slowly. Recoded into Java and Processing, and using a streamlined "wire-and-dots" visual representation, the model achieved acceptable run times (near real time) and a capability to generate files that can be rendered, not in real time, into fairly realistic movies. The subway model is primarily built in NetLogo, as the smaller crowds in the subway station make lesser demands on processing power

AGENT-BASED SIMULATION

The essential idea of agent-based modeling and simulation[1] is that many phenomena, even very complex ones, can best be understood as systems of autonomous or semi-autonomous agents that are relatively simple and follow relatively simple rules for movement and interaction. The fundamental feature of an agent is the capability to make independent decisions: an agent is a discrete individual with a set of characteristics and rules governing its behaviors and decision-making capability. An agent:

- may have additional second-level rules that modify its first-level rules of behavior;
- can function independently in its environment and in its dealings with other agents, at least over a limited range of situations;
- is goal-directed, having goals to achieve (not necessarily objectives to maximize) with respect to its behaviors;
- is flexible, and has the ability to learn and adapt its behaviors over time based on experience, which generally requires some form of storing and retrieving information at the agent level.

Thus, by specifying a few fairly simple, intuitive rules of behavior for agents and a fairly simple set of rules by which they interact, the modeler can represent complicated collective and system behavior, such as mass egress from a large structure, in a way that replicates reasonably well the variety of things people will actually do in such situations. A model can easily include a number of different types of agents, each with different sets of rules of behavior. It is no longer necessary to assume away much of the variety and changeability that characterize real-life activity.

A number of agent-based simulation software packages have been released within the past four years. This new software, combined with continuing increases in computing power, can handle much larger numbers of individual agents than the best discrete-event packages a few years ago. For example, the stadium model developed for this task can represent, in real time, a crowd of 70,000, on a current-generation laptop computer. This is more than an order-of-magnitude improvement in both scale and speed over the discrete-event simulation available five years ago.

New software toolkits also use visual input via a graphical user interface (drag-and-drop icons and pull-down menus) and produce visual, easily understood, movie-like output. The user can change variable inputs in the interface window quickly, in some cases even in mid-run, via point-and-click. This capability makes "what-if" and sensitivity analyses much easier and quicker; it also enables non-specialists to interpret model outputs and obtain analyses of alternative actions that authorities or incident managers could take.

---

[1] There are few survey articles about agent-based simulation so far. An especially good one, from the standpoint of explaining the approach and its scientific principles, is Bonabeau, 2002, "Agent-based modeling: Methods and Techniques for Simulating Human Systems," *Proceedings of the National Academy of Sciences* 99(3): 7280-7287. In terms of tools and techniques, however, this article is already somewhat dated because of the proliferation of new software and the resulting variety of new applications. A recent, less technical overview, focusing on the new software, is Samuelson and Macal, "Agent-Based Simulation Comes of Age," OR/MS Today, August 2006, 34-38. Both articles contain, in turn, numerous references to other useful sources.

CROWD DYNAMICS

The HSI team found some useful prior work on crowd dynamics. In particular, Helbing et. al.[2,3] constructed models based on the physics of granular flows for materials such as sand or salt. Granular substances often form a "cap" or arch-like structure that inhibits flow toward a drain. An intriguing and somewhat counter-intuitive finding from Helbing's agent-based simulation is that placing an obstacle, such as a two- to three-foot-tall ashtray, in front of an entrance to a stairway or escalator may facilitate smoother movement, as the obstacle prevents the people from forming a "cap" structure. He went on to state that this measure has shown some benefit in controlled live tests.

Additional searching reinforced the HSI team's impression that the relevant literature is sparse, especially in the U.S. One book, Schreckenberg et. al.,[4] seemed to summarize the state of current research fairly well — again using almost entirely European sources and examples, with no explosion-induced instances. In particular, the chapter by Deborah Withington extensively analyzed sound-based direction. Her conclusion that such direction is superior to vision-based direction methods, such as lighted "Exit" signs and lighted arrows pointing to nearby exits, motivated the HSI team's conclusion that direction of egress is important and under-analyzed, especially in cases when egress can be directed and re-directed in response to the current emergency situation.

The U.S. source that appeared to have done the most relevant prior work was Redfish Group, in Santa Fe, New Mexico.[5] Their agent-based models of evacuation from the valley around Santa Fe in the event of wildfire (still a work in progress when this project began) and of likely evacuation problems following an explosion at Santa Fe's annual "Zozobra" festival seemed most relevant to the study at hand. Accordingly, HSI engaged Redfish Group as a subcontractor.

In the Zozobra study, Redfish Group found that both additional exit routes and clear direction to alternative exit routes seem effective in greatly reducing panic and trampling. In the wildfire study, they found that pre-incident planning of alternative responses and exit routes is critical, as some exit routes may be unsafe depending on the direction of spread of the fire, and there is little time and limited capability to inform people of preferable routes.

A number of recent analyses of the 2001 attack on the World Trade Center identified limitations of stairway capacity as a possible contributor to additional deaths and injuries. Also, some of these analyses, in particular Chertkoff and Kushigan[6], noted the particular importance of conflict between people trying to leave and emergency responders trying to enter the upper floors: the building designers apparently did not take emergency responders' ingress into account in deciding how much stairway capacity would be needed. The World Trade Center has also been the focus of the most comprehensive assessment of long-term effects.[7] Prolonged exposure to the airborne debris plume appears to have had major consequences of an intensity and a scope previously

---

[2] http://angel.elte.hu/~panic/ (Helbing Web site)

[3] D. Helbing et. al, Simulating dynamical features of escape panic. *Nature* 407, 487-490 (2000).

[4] Schreckenberg, M., & Sharma, S. D., *Pedestrian and Evacuation Dynamics*, Springer, 2001.

[5] www.redfish.org

[6] Chertkoff, Jerome M. and Russell H. Kushigian, *Don't Panic: the Psychology of Emergency Egress and Ingress*, Praeger, 1999.

[7] http://www.nyc.gov/html/doh/html/wtc/about.html wtc health registry

unsuspected, a finding which further underscores the importance of egress following an attack. This result also suggested to the HSI team, independent of the modeling, the potential value of designing in ingress routes for the exclusive use of emergency responders, such as closed-access routes under the seats in the stadium.

Chertkoff and Kushigan[8] studied a number of disasters in which crowd dynamics played a crucial role, and identified ten factors that determine severity of consequences. These factors combine in such a way that, when several of the factors occur together, the combined effect is much greater than the sum of individual effects. Among the most important factors are extreme crowd densities, a perceived imminent threat to safety, poor or absent direction of egress, and congestion of the areas beyond obvious choke points.

While the modeling work was in progress – specifically, in late August 2006, after the stadium model was mostly completed and the subway model was partially coded – new work on representing crowd dynamics became available.[9] This "Continuum Crowds" algorithm, in contrast to previous approaches, enables agent-based representations of crowds to yield realistic representation of crowd dynamics as documented in the literature. Hence the HSI team judged this new method more accurate than previous approaches, and adopted it for the subway model.

SELECTING AND REPRESENTING VENUES

HSI and Redfish Group considered a number of venues as candidates for the stadium model. Of several seemingly appropriate venues, PNC Stadium, the major-league baseball stadium in Pittsburgh, stood out because the city and the stadium authority had conducted a live egress exercise there within the past few years. This meant that limited actual data were available against which to check modeled rates of flow through various passageways and choke points in the structure. Redfish Group analysts traveled to PNC Park and studied its features extensively, observing crowd movements before, during and after two baseball games, and corresponded with the chief of security for the facility.

The HSI team met with the chief of security for the Washington Metropolitan Area Transit Authority (WMATA), to learn of issues affecting subway stations in the event of an attack. He mentioned that, in its own studies, WMATA had concluded that one of the most useful actions it could take was to upgrade the public address system's speakers: the system's ability to process information and get it to its own staff was deemed good, but the announcements to the public were often largely unintelligible.

REPRESENTING EXPLOSIONS

The team consulted well-known reference sources both for the characteristics of explosions[10] and the effects of various shock- and shrapnel-absorbing materials.[11] Explosions harm people in three ways: the direct shock wave and, in enclosed areas, reflections of the shock wave; overpressure; and flying shrapnel.

---

[8] Chertkoff, Jerome M. and Russell H. Kushigian, *op.cit.*

[9] A. Treuille, S. Cooper, Z. Popovi, "Continuum Crowds," ACM Transactions on Graphics 25(3) (SIGGRAPH 2006) — see also http://www.washington.edu/projects/crowd-flows/

[10] Kinney, Gilbert F., and Kenneth J. Graham. *Explosive Shocks in Air*, 1985.

[11] Biggs, John M., *Introduction to Structural Dynamics*, McGraw-Hill, 1964.

For the stadium model, the team judged it sufficient to represent explosions simply as a sphere, within which everyone would be killed or incapacitated, and outside which everyone would be unharmed. The team adjusted the radius of the sphere to yield numbers of casualties consistent with reported suicide vest explosions elsewhere: typically 20 to 60 per explosion, in the moderately dense crowds in which such bombs have maximum effect. This approach results in what is probably a somewhat overstated effect in the denser crowds in the stadium, but it is not a large enough difference to change the findings substantially.

In the stadium, this representation is fairly accurate, as most of the casualties are caused by shrapnel. The range of shrapnel exceeds that of the shock from the blast. Overpressure is negligible, except for explosions inside the large rotunda stairways, because the shock wave simply continues upward into the open air. Even in the stairways, the team concluded, based on the equations in the Kinney and Graham book, that the presence of windows that would probably blow out most likely rendered overpressure unimportant. This simple representation saved substantial computer capacity, which was strained to its limits in this model because of the large number of people potentially present.

In the subway, however, the team deemed that more accurate representation of explosion effects was desirable. Hence the team constructed an agent-based representation of the explosion. Shock waves are modeled as agents radiating in all directions from the explosion's core and reflecting off hard surfaces, and bits of shrapnel are (small) agents that radiate linearly, with exponentially decaying velocity, from the explosion's core. People-agents accumulate impacts of direct and reflected shock waves and shrapnel and use probability tables to determine whether they die, become severely incapacitated, or suffer lesser injury. This representation follows the equations in the reference and produces results that appear quite accurate.

AGENT-BASED MODELS DEVELOPED FOR THIS TASK

HSI developed two models for this task: one of a stadium and one of a subway station. The process involved considerable learning as the project progressed, and additional learning and modification continue.

These models were designed to provide generic testbeds to evaluate improvised explosive device (IED) mitigation strategies through the employment of one or more selected technologies. They represent the likely movement of large numbers of people in the selected venue before and after one or more detonations of IEDs, assumed to be around the size and force of a suicide vest consisting of about 25 pounds of explosive (TNT), surrounded by shrapnel.

Stadium

The first model HSI undertook was the larger and more complex. This was partly because the stadium was the scenario of greatest interest to the sponsor and partly because tackling this task first enabled HSI to learn more quickly how difficult the technical challenges would be.

Although Redfish Group began building the model in NetLogo, the team soon realized that execution was too slow to be useful. Therefore, the team switched to Processing, an add-on to Java, for this model. Although Processing is harder to learn and less well documented than NetLogo, it is also available in open source and is much more efficient. This made it possible to represent the behavior of as many as 70,000 people realistically, in real time. (This is more than the stadium would typically hold for a baseball game, as about 40,000 would be a good-sized crowd.)

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.

The output includes this 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.

Countermeasures of interest were permitting egress via the field and improved means to direct evacuation. In the workshop mentioned above, bomb squad commanders noted that they obtain much better compliance with orders to evacuate when they can tell people which way to go. In most venues, the primary source of such information is lighted exit signs; these are typically red, and placed high enough that people can easily see them above other people's heads. The problem with this is that smoke rises, and fire tends to generate mostly red-orange-yellow light that reflects in various directions. Lights in the blue-green range, on the floor, have been found to work better in airplane cabins. In addition, smoke irritates people's eyes, making vision less reliable in general. Therefore, sound-based systems may be preferable.

Figures 1 and 2, below, are screen shots taken at a selected time for two runs, one with the crowd displaying typical confusion and the other with improved guidance to exits. These runs used a crowd of 40,000 people. In both cases, egress onto the field is permitted. In each case, three suicide vests have been detonated at 20-second intervals: one in the middle of the stands behind third base, directly facing the viewer; a second in the middle of the right-field stands; and a third behind home plate (home plate is at the top left corner of the field in this view.) The blast areas can be (barely) seen as small red circles. The difference in patterns of crowd movement, however, is plainly visible.

The graph above the picture of the stadium shows numbers of casualties from the explosions and from trampling, and total number of people who have left the stadium area. (People are not counted as "exited" until they leave the vicinity, not just when they get outside the building.)

Things to notice:

- Guidance to exits greatly expedites egress but does not reduce tramplings.
- The large rotunda stairways (behind the home-plate stands and behind the left-field foul pole, i.e., just past the corners away from the viewer) still have few people in them five minutes into the scenario. This means rapidly sequenced attacks would have more effect if all of them were in the stands, rather than if some of them were in the stairways.

- The areas affected by bomb blasts are small relative to the venue, so they have little effect on crowd movement other than the general tendency to try to leave.

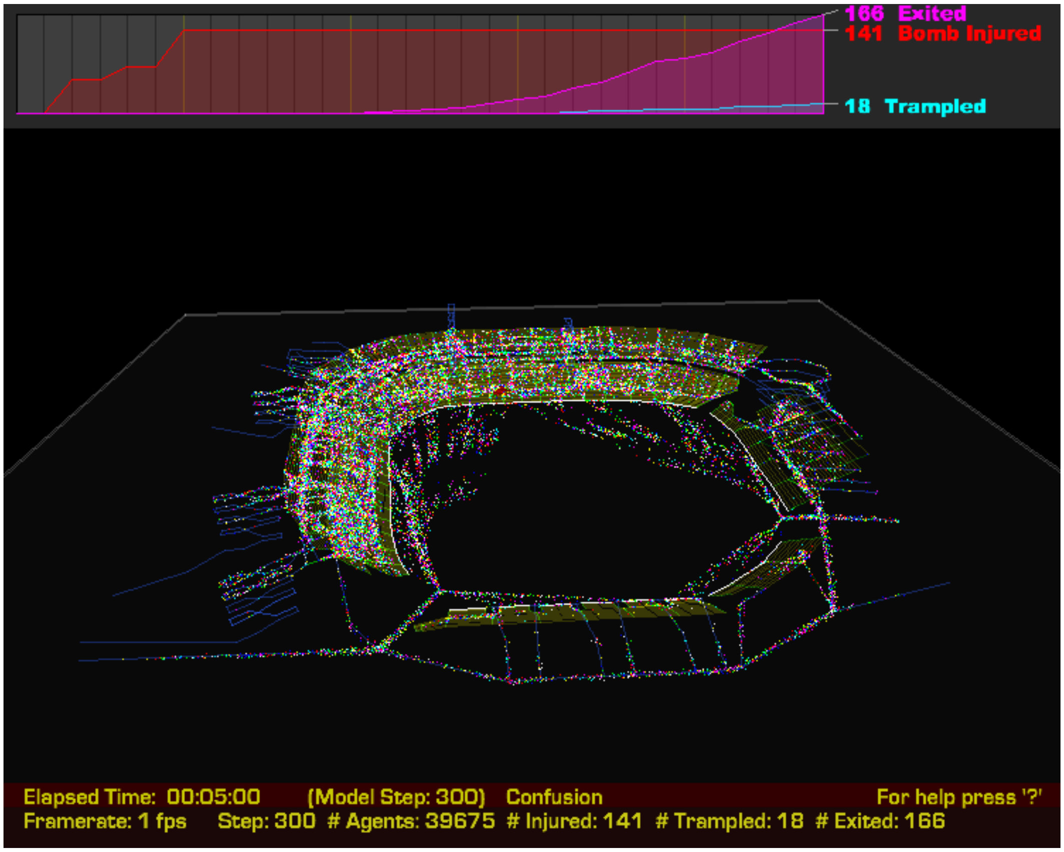

```
Elapsed Time:  00:05:00     (Model Step: 300)   Confusion                          For help press '?'
Framerate: 1 fps     Step: 300  # Agents: 39675  # Injured: 141  # Trampled: 18  # Exited: 166
```

Figure 1: Stadium model screen shot at 5:00 simulated time, following three IED detonations — crowd dynamics with confusion (no guidance to exits), 40,000 people

NOTE: multi-colored dots are people (the colors have no meaning), green background patches are unoccupied seats, and the lines of egress are shown in dark blue. The graph at the top is a multi-variable time series chart, with number of bomb casualties, number exited, and number trampled on the vertical axis.

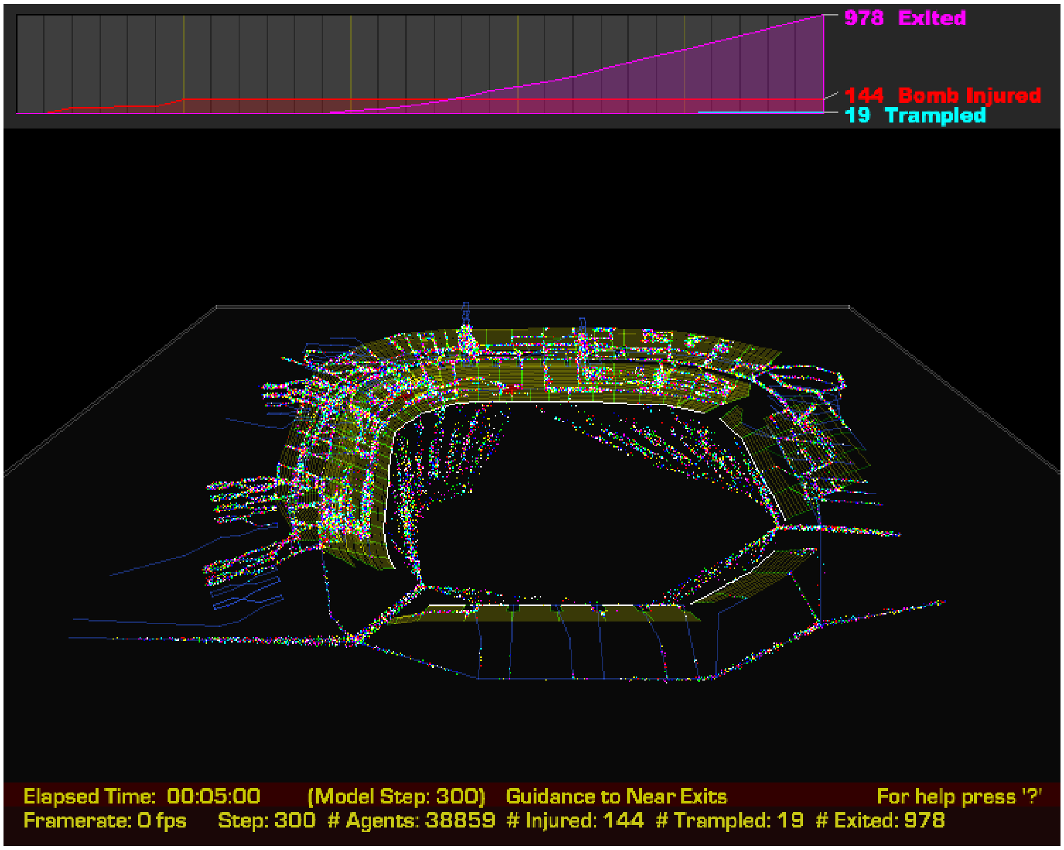

Elapsed Time: 00:05:00 (Model Step: 300) Guidance to Near Exits For help press '?'
Framerate: 0 fps Step: 300 # Agents: 38859 # Injured: 144 # Trampled: 19 # Exited: 978

Figure 2: Stadium model screen shot at 5:00 simulated time, following three IED detonations — crowd dynamics with guidance to exits, 40,000 people

Subway

The subway model is based on the Metro Center subway station in Washington, D.C. The model uses only the upper of the two levels of the actual station. The HSI team chose this particular station because data were readily available and because it is larger and more complicated than most subway stations (Metro Center is a transfer station between two subway lines.) Thus, while the current version of the model uses the simplified representation of a more typical station, and only one set of trains, it could readily be expanded to contemplate scenarios involving multiple-line stations.

Note that the depiction is of two levels: the upper area on the screen is the lower level in the station, where the trains run, and below that on the screen are the two upper-level concourses. The escalators at the innermost ends of the concourses connect to the escalators on the platforms. (This takes some getting used to, especially for viewers unfamiliar with the actual subway station that served as the example.)

This 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.

People are represented as agents. Locations are represented as patches, and each person-agent 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.

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.

The following figure is a screen shot of the subway model, for illustrative purposes – it is not readable at the size that fits into this page. We analyzed screen shots like this taken at selected times for two runs, one with no countermeasures in use and the other with baffles to absorb shock and shrapnel spaced along the platforms. These runs used a crowd of 1,000 people. In each case, a suicide vest was detonated on the near side of the lower level platform 15 seconds into the run, and another was detonated in the upper right concourse 45 seconds into the run. The graphs below the depiction of the station show numbers of casualties from the explosions and from trampling, and number of people remaining in the station.

From runs of this model, we observed:

- Baffles protect people directly on the other side of the baffle from the explosion, but it is easy for the attacker to select a place of detonation that minimizes the beneficial effect.
- Bomb casualties are shown as red X’s, tramplings as purple X’s. There are two tramplings in the “no-technologies” run and none in the “baffles” run; the team attributes this difference to random variation. The two trampled people can be seen, with some difficulty, in the right upper concourse.
- The station empties out fairly quickly in either case.

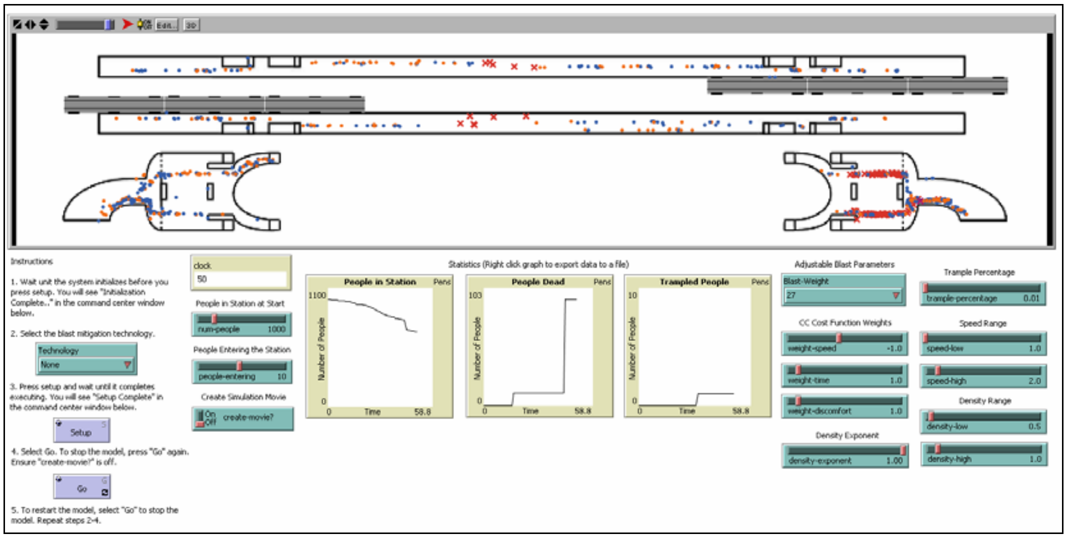

Instructions

1. Wait unit the system initializes before you press setup. You will see "Initialization Complete.." in the command center window below.

2. Select the blast mitigation technology.

Technology: None

3. Press setup and wait until it completes executing. You will see "Setup Complete" in the command center window below.

4. Select Go. To stop the model, press "Go" again. Ensure "create-movie?" is off.

5. To restart the model, select "Go" to stop the model. Repeat steps 2-4.

clock: 50

People in Station at Start: num-people 1000

People Entering the Station: people-entering 10

Create Simulation Movie: create-movie? Off

Adjustable Blast Parameters: Blast-Weight 27

CC Cost Function Weights: weight-speed -1.0; weight-time 1.0; weight-discomfort 1.0

Density Exponent: density-exponent 1.00

Trample Percentage: trample-percentage 0.01

Speed Range: speed-low 1.0; speed-high 2.0

Density Range: density-low 0.5; density-high 1.0

Figure 3: Subway model screen shot at 0:50 simulated time, following second of two IED detonations — no technologies used

FINDINGS FROM THE MODELS

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 clearly beneficial 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.

SUMMARY AND CONCLUSIONS

The HSI team's agent-based simulation models of a stadium and a subway station provided useful and interesting insights into how one or more explosions would affect mass egress and the eventual number injured or killed. HSI's analysis supports the conclusions that this modeling approach is valuable for assessing the likely benefits of proposed mitigation technologies, including improved information systems, and that such models would also be useful for assessing the costs and inconvenience other technol-ogies, such as screening, would impose. From this study, agent-based simulation appears to provide greater realism, adaptability, ease of use and ease of interpretation than traditional simulation methods. It also uses a visual, intuitive interface and produces movie-like output, making these models much easier to use, even by non-modelers, to conduct "what-if" analyses. Therefore, such models appear valuable for pre-incident planning and exercises, as well as the technology assessment purpose for which they were designed. In particular, they indicate the potential value of information systems to guide egress.

The sponsor considered some of the other resulting policy recommendations interesting enough to restrict dissemination, marking them "For Official Use Only." The lead author may discuss some of these topics on request, on a case by case basis.

The HSI team strongly recommends agent-based modeling as a way to assess benefits of proposed technologies, to assess the likely impacts of policy options, and to support planning and exercises. The HSI team therefore also strongly recommends further investigation of the uses of simulation in preparedness for incident management, particularly assessing the likely effects of information systems.