Proceedings of the 2008 Winter Simulation Conference
William E. Weiss

DYNAMIC SECURITY: AN AGENT-BASED MODEL FOR AIRPORT DEFENSE

William E. Weiss

The MITRE Corporation
7515 Colshire Drive
McLean, Virginia 22102, USA

ABSTRACT

The Department of Homeland Security (DHS) shifted the focus of airport security in 2004 to incorporate the need to continuously and rapidly adapt security to shifting threats. MITRE is developing a Dynamic Security Airport Simulation as part of a MITRE-sponsored research project in which attacker and defense behavior in the airport environment are modeled. The simulation accepts threat vectors (path-weapon combinations) from other software or the user and models the performance of the airport defense against those threat vectors. The simulation includes two intelligent agents: the attacker and the defense. These agents model the behavior of those two entities; their logic includes both decision making and learning.

1 INTRODUCTION

The Transportation Security Administration, a branch of the Department of Homeland Security (DHS), is responsible for U.S. airport security in conjunction with airport operators. DHS espouses a layered, adaptive security concept similar to that used in cyber security (U.S. Department of Homeland Security 2004). The objective of this type of security is to adapt security measures in proportion to changing threats. Along with outright attacks, U.S. airport operators can expect probes from attackers, designed to test their defenses. These probes supply the attacker with information on the airport's defenses, but can also supply the airport's defense with information on the attacker if the probes are detected. Although these probes occur at a much lower rate than probes of cyber security, airport operators must still be prepared to act on information obtained from them.

In 2004, DHS shifted the focus of airport security, incorporating the need to continuously and rapidly adapt security to shifting threats (Chertoff 2005). A DHS strategic review that followed also emphasized this strategy, along with the need for analytic tools to help match security to perceived threats.

At this time there are few tools available to the airport security coordinator to test how well the airport is safeguarded against changing threats. Only a very small number of attack scenarios can feasibly be tested in live exercises. Red Teams are somewhat less expensive and difficult to set up than live exercises, but still cannot be used to assess large numbers of threat scenarios.

Thus MITRE, a not-for-profit organization chartered to work in the public interest, began to investigate airport security risk as a function of perceived threats, measuring how well security designs and procedures match up against dynamic threats. One aspect of that research involved assessing the threat vectors (path-weapon combinations) most likely to result given at least some knowledge of attacker's goals and capabilities. The threat vectors from that research, or threat vectors generated by the user, feed the Dynamic Security Airport Simulation.

2 SIMULATION DESIGN

The simulation is being developed using ExtendSim Version 7, which lends itself to agent-based simulation. Modelers build simulations with ExtendSim by adding icons (called "blocks") to a worksheet. Each type of block has a different function, and ExtendSim comes with many dozens of blocks. Any set of blocks can be combined into a hierarchical block, or h-block. Although many simulations can be written using the included blocks, the user can write his own code to perform specialized functions. (This is likely to be necessary in more-complex simulations, although ExtendSim supplies shortcuts to minimize the amount of code required.) The blocks are connected to define the network and provide pathways for both data and the simulated items (people) to traverse. The simulation is then executed. Included two-dimensional graphics illustrate model flow and help with debugging.

The complex, adaptive system being simulated is one of passengers progressing through an airport terminal, en route to their departing flights, and of airport and airline employees en route to their work areas. The model assumes that some of these people are attackers; each at-

tacker's decisions and actions are modeled via an attacker agent, while the airport's defense decisions and actions are modeled via a defense agent. The operation of the model also assumes that the defense obtains and acts on information about people as they progress through the simulation.

The simulation's primary process is the series of encounters an attacker has with the sensors and barriers established in the simulated airport by the defense. The attacker reacts to them and his actions are modified as a result. The defense, in turn, reacts to the attacker's actions depending on the defense's rules of engagement and what its sensors detected.

The model includes many measures of the attackers' and defense's success. In the most basic sense, the attackers' success is measured by the number that reach the aircraft, while the defense's success is measured by the percentage of attackers apprehended. Other measures, such as the number attackers who turn away and do not pursue their intended path as a result of questioning, are also tracked, as are the locations within the airport at which attackers are detained.

3 DEFENSE

The defense is composed of a network of sensors and barriers, coupled with logic for reacting to attacker attributes and actions.

3.1 Defense Structure

Each h-block depicted in Figure 1, such as Ticket Counter 2, contains the logic for the attacker-defense encounter that may occur at that location. This model is not designed to be a detailed simulation of passenger movements through the airport terminal; thus, only those aspects that relate to airport security are modeled in the defense structure.

3.2 Defense Agent

For barriers, the logic for each defense agent is contained in the Barrier and Defense Action blocks, described in the Barrier section below. For sensors, the agent logic is in the Sensor and Defense Action blocks, shown in Figure 2.

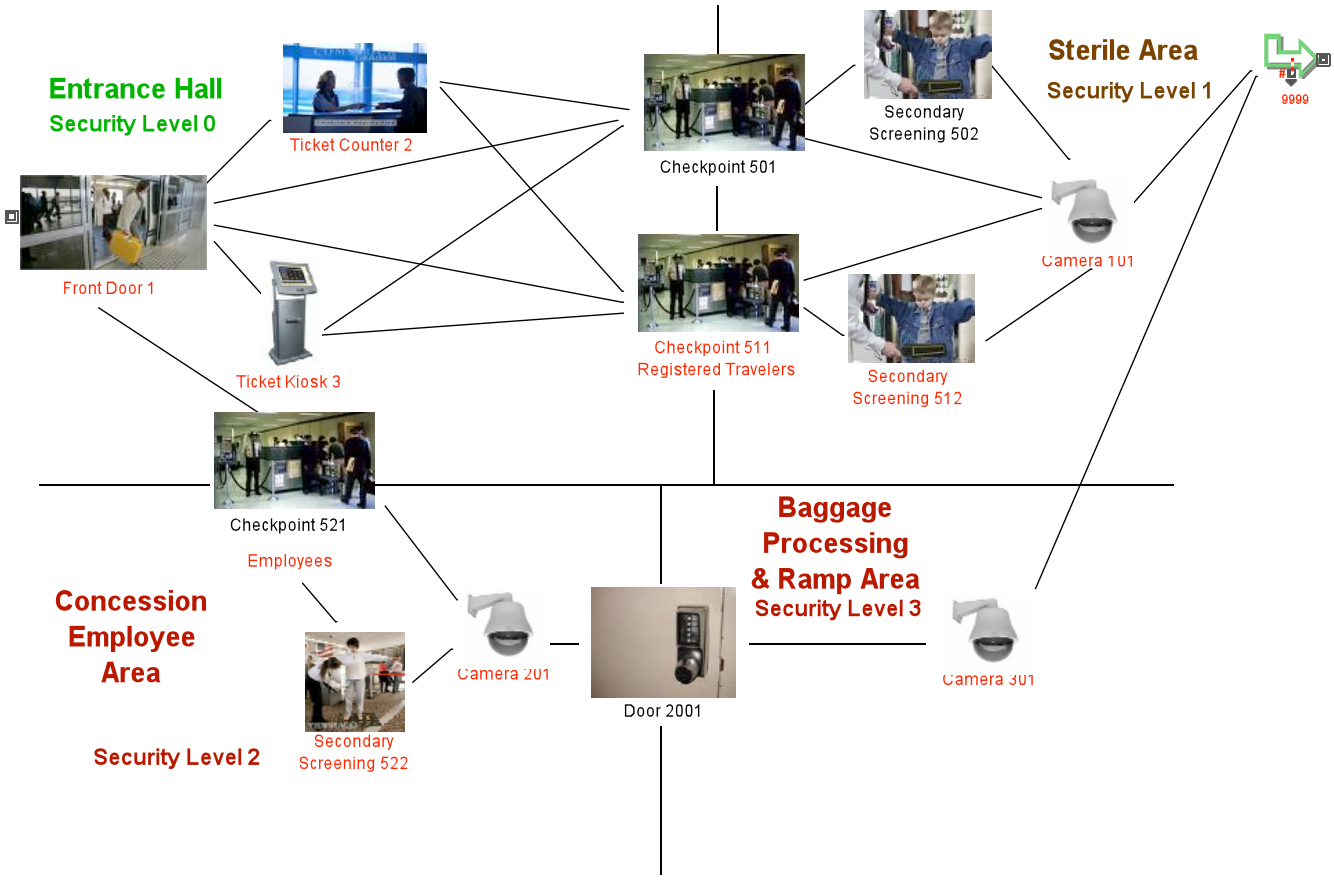

Figure 1: Airport Defense Network Layout

The high-level logic for a sensor is shown in Figure 2, a ticket counter and typical sensor location. In this case the sensor is human observation in the form of an airline employee, a monitored video camera, or a security officer stationed near the counter.

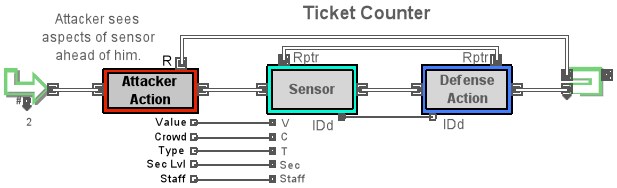

Figure 2: Typical Sensor Layout

This sensor (human eyes) has a defined probability of detecting one of the following:

- Nothing
- Weapon type 1
- Weapon type 2
- Weapon type 3
- Suspicious behavior
- Person of interest

Each weapon type could be any kind of weapon specified by the user or even a characteristic of a person. For example, the human eyes at the ticket counter may notice that a passenger has symptoms of a possible contagion. If the existence of that condition is considered a "weapon" by the defense in the simulation, there is a probability that the passenger could be questioned by the defense.

Note that the identification of a weapon, person of interest, or behavior is separate from the actual existence of each (within the simulation). For example, a person could be mistakenly identified as possessing a weapon, as in real life.

3.3 Defense Decisions

Once a person has been sensed by the defense, the defense decides what actions to take regarding that person. If one of the above "weapons" (each weapon type, suspicious behavior, and persons of interest are all treated as "weapons" by the model) is identified (rightly or wrongly), that information is passed on to the Defense Action block, where one of several actions is taken:

- No action
- Sensing immediately repeated (as in a person's second pass through the walk-through metal detector after removing additional articles)
- Secondary screening inserted into path (essentially, hand-wanding at the airport security checkpoint)
- Hand searched
- Questioned

If a person is hand searched or questioned, there is a probability that he may be detained. Alternatively, the defense's scrutiny of that person may be increased.

3.4 Barriers

Barriers are similar to sensors in the simulation (as shown in Figure 3), but with the additional function of impeding the progress of the person attempting to pass through. Barriers require a "key" for passage, which in the real world could be an actual key, a pass code, a radio-frequency-ID-equipped identity card, or similar. The simulation recognizes three different versions of a key—which may be a valid key, a forged key, or a stolen key, specified at the user's discretion—all of which have different, user-defined probabilities of opening the barrier. And, a different key is required to enter each of three levels of security.

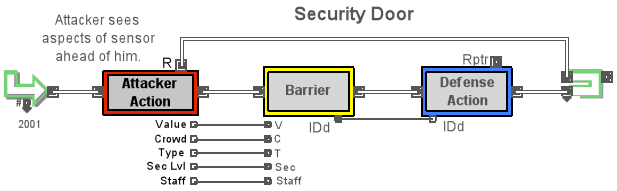

Figure 3: Barrier Layout

Repeated, failed attempts to open a barrier may result in the defense increasing its scrutiny of that person.

Barriers have a greater number of potential outcomes than sensors. A person (referred to as "pax" in Figure 4, below) may open a barrier and continue; fail to open a barrier and retreat; may already be in retreat and may simply pass by a one-directional barrier; or may be in retreat, fail to open a bidirectional barrier, and be caught behind it by the defense.

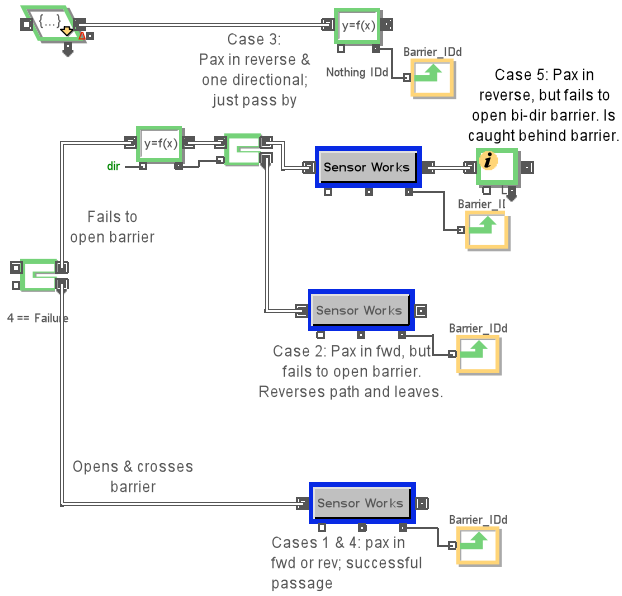

Figure 4. Barrier Detail

3.5 Defense Characteristics

Two characteristics affect the defense's actions: scrutiny and staffing.

3.5.1 Scrutiny

Scrutiny is the level of attention the defense pays to each person in the simulation. Essentially, it is the ability of the defense to track a person or item of interest and focus resources on him. Although scrutiny affects the defense's actions, it is associated with each simulated person.

Each item simulated within an ExtendSim simulation can carry attributes with it; these can be either numeric or text. Attributes such as "Has Weapon Type 1" describe the current state of a simulated person, while attributes such as "IDd Weapon Type 1" describe what the defense believes to be true about a simulated person.

Scrutiny is a similar attribute and can be one of three values:

- Ignore
- Increased
- Intense

A person enters the simulation with a scrutiny level of "Ignore." However, sensing results can increase the defense's scrutiny of a person. The presence of a higher level of scrutiny implies that more attention is being paid to a person and thus, in the simulation, increases the chances that the defense will identify a weapon, suspicious behavior, or a person of interest.

3.5.2 Staffing

In the simulation, as in the real world, staffing affects the probability of identifying a weapon, suspicious behavior, or person of interest. It represents the degree of attention paid to the sensors. Staffing is specified for each sensor at one of the following levels:

- None
- Low
- Nominal
- Full
- Over (overstaffed)

If a sensor in the simulation is unstaffed, there is zero probability of sensing anything. (If one were modeling an unstaffed but fully automatic sensor, then one would specify a nominal staffing setting for that sensor.) Low staffing implies a reduced probability of sensing anything (versus nominal) and full staffing implies an increased probability of sensing anything. Finally, overstaffing implies that one hundred percent of weapons, suspicious behavior, and persons of interest will be identified. (The probabilities of all of these characteristics are specified in the model's built-in database and can be modified by the user.)

4 ATTACKER

A primary concept in the simulation design is that of an intelligent attacker that can discern what lies ahead of him and will react accordingly.

4.1 Attacker Agent

The attacker agent is embodied in the Attacker Action block, the contents of which are shown in Figure 5. Although any of several different actions may take place in that block, the simulated attacker acts based on what he discerns:

- Monetary value of the environment
- Crowd size
- Type of sensor or barrier confronting the attacker
- Staffing level
- Security level beyond the sensor or barrier

4.2 Attacker Characteristics

Each attacker has several characteristics, one of which is weapons. An attacker may have one of the weapons described in Section 3.2 above, or the defense, via one of its

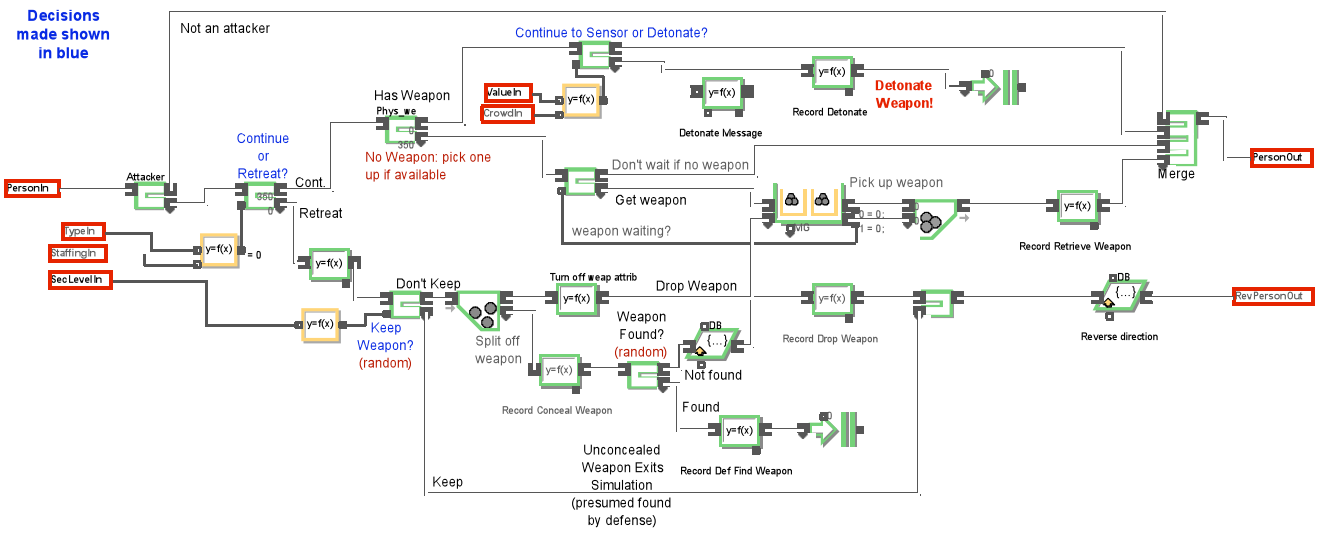

Figure 5: Attacker Agent Logic

sensors, may have identified the attacker as having one of these weapons.

Attackers also have a direction; this direction specifies whether the attacker is attempting to move forward along the path that he entered the simulation with, or if he is retreating by attempting to move backward along that path.

A final, important attacker attribute is scrutiny, described above.

4.3 Attacker Decisions

Figure 5 illustrates (in blue) the decisions that an attacker must make when confronted with a sensor or barrier. (Note that, if a sensor exists in the simulation but cannot be discerned by the attacker, the Attacker Action block is omitted and the attacker makes no decisions.)

First, an attacker must decide whether to advance or retreat. He uses the type of sensor in front of him and its staffing level as factors in that decision.

If the attacker elects to retreat, he may drop his weapon to minimize his chances of detection. The security level of the surrounding area influences this decision. And if he drops his weapon, he may decide to hide it for a future attacker to retrieve, depending on his rules of engagement.

If the attacker decides to continue but has no weapon, he may pick up a weapon if one has been left and hidden by a previous attacker.

If the attacker decides not to retreat and has a weapon, he then decides whether or not to deploy his weapon on the spot. Factors in this decision include the value of the surrounding area and the crowd size. This decision on the attacker's part reflects his propensity for satisficing (Simon 2001).

If the attacker elects to continue, he advances to the next sensor or barrier in his path.

5 AGENT-BASED DECISION COMPLEXITY

The agents in this simulation for the most part respond to routine events using base-level rules (North and Macal 2007). Learning is present within the simulation via the staffing and scrutiny parameters, and is present externally as model is incrementally used to improve the airport defenses.

Because the model is a work in progress, more-sophisticated decision making will be added as development continues. An obvious area for more sophistication is will be to add rules that will change the existing base-level rules as the modeled scenario develops—in essence, adding more-advanced coordination and learning. Some sample areas include the following:

- Coordination between attackers that would allow parts of weapons to be brought into the airport for later assembly.
- More-advanced learning that would alter attacker strategies based on defense actions taken in response to earlier attackers.
- Learning that would enable the defense to alter its defensive scheme in response to discerned attack patterns.

Obviously, there are other opportunities in the model for more-sophisticated decision making as well. The modular structure of the model and its data and internal communication mechanisms should support this level of sophistication.

6 RUNNING THE SIMULATION

The simulation is run by first combining defense-related h-blocks into a network; duplicate copies of the blocks can be made to ease constructing the network. Although it makes intuitive sense to place the defensive h-blocks in positions relative to their real-life positions, in fact the placement of the h-blocks on the worksheet is unimportant. (This is because ExtendSim's "catch and throw" capability is used for data transfer between defensive h-blocks for flexibility and to obviate the need for a network of connection lines at that level. However, within those h-blocks, people and data move along connection lines, as shown in Figure 5.) Thus, an important step in specifying the network is ensuring that each defense h-block is "aware" of the defense blocks that precede and follow it. Each Catch and Throw block is identified by a unique value to ensure that simulated passengers and employees follow their correct paths.

6.1 Inputs and Outputs

Input data is specified in two places: a Microsoft Excel file that holds frequently-changed data, and the model's internal database, supplied with ExtendSim, which holds the simulation's parameters.

The Excel file holds a set of paths that each type of person in the simulation follows. Although any number of person-types can be defined, presently, there are three:

- Regular passenger
- Registered traveler program participants
- Airport and airline employees

The Excel file also holds both detailed simulation outputs, useful for debugging, and outputs such as the number of attackers detained in each simulation run. To facilitate using the model, its basic outputs are shown in the main simulation worksheet. Outputs can also be shown on the simulation Notebook, a customizable window supplied by ExtendSim that is useful for viewing input data and results.

The model's internal database contains all of the parameters required for every aspect of the simulation. Some examples include the probabilities of detecting each weapon type by each sensor type, the probabilities of each type of key opening a barrier, and the probabilities of attackers making certain decisions. The database has a convenient graphical interface useful for viewing and modifying simulation parameters.

6.2 Executing the Simulation

The simulation can be executed either with or without graphics. ExtendSim has both 2D and 3D graphics; the 2D graphics are useful for debugging the simulation. The 3D graphics, while available, have not been set up for this simulation. The simulation executes, with graphics off, in about one second for 1,000 passengers.

REFERENCES

Chertoff, M. *Statement of DHS Secretary Michael Chertoff before the United States Senate Committee on Homeland Security and Governmental Affairs: "Second Stage Review"*. 7/14/2005. Electronic Document: <http://hsgac.senate.gov/_files/071405Chertoff.pdf>

ExtendSim User Guide Release 7. 2007. Imagine That, Inc., San Jose, California, USA. <http://www.extendsim.com/>

North, M. and C. Macal. 2007. *Managing Business Complexity*. New York, NY. Oxford University Press.

Simon, H. 2001. *The sciences of the artificial*, Cambridge, MA: MIT Press.

U.S. Department of Homeland Security. *Securing Our Homeland*: *U.S. Department of Homeland Security Strategic Plan*. February 2004.

AUTHOR BIOGRAPHY

William E. Weiss is a lead engineer in MITRE's Center for Advanced Aviation System Development. He has written numerous aviation-related simulations over the course of a 28-year career. In 1989 he co-authored the National Airspace System Performance Analysis Capability (NASPAC) model, the first simulation of all controlled air traffic over the continental U.S. NASPAC is still being used by the Federal Aviation Administration to evaluate technical and policy changes on national airspace system operations.