Using Simulation to Produce a Data Fusion Decision Support Tool for the Assessment of Manmade and Natural Disasters

Donald E. Brown, Justin R. Stile & Nicole M. White
Department of Systems and Information Engineering, University of Virginia
151 Engineer's Way, P.O. Box 400747, Charlottesville, VA, 22904-4747
Phone: (434) 924-5393, Fax: (434) 982-2972
Email: Brown@virginia.edu, jrs5p@virginia.edu, nmw5h@virginia.edu
Contact Author: Justin Stile, jrs5p@virginia.edu

Abstract: *Effective performance in complex, dynamic command and control environments depends on high quality awareness. In such situations, an Emergency Command Center's (ECC) ability to deal with multiple disparate sources of information is complicated by the volume of information available from human experts and automated decision support tools. In the event of a natural disaster, large amounts of information, both accurate and inaccurate, must be assessed by the ECC. The data to be analyzed is reported to the ECC by field agents, or sensors. These sensors are disbursed throughout the disaster area and vary from trained professionals to civilians. The correct assessment of the information reported by these sensors is necessary to support a variety of authoritative decisions made by the ECC. This paper focuses on the creation of the intelligent agent based reporting simulation that produces the ground truth and field reports necessary for data fusion to analyze.*

Key Words: Simulation, Agent, Data Fusion, Natural Disaster, Sensor

1 Introduction

Effective performance in complex, dynamic command and control environments depends on high quality awareness. In such domains, high levels of situation awareness rely on detecting the important elements within the areas of concern. In such situations, an Emergency Command Center's (ECC) ability to deal with multiple disparate sources of information is complicated by the sheer volume of information available from both human experts and automated decision support tools. In many cases, information fusion technologies can improve situation awareness through the mining and integration of data from multiple disparate sources. Such knowledge-based systems can represent the current state of complex environments so that inferences can be drawn about current and potential future events. By exploiting and compressing sensor systems and human reporting, information fusion can assist in forming more effective assessments, and yielding improved decision-making. However, with volumes of information, generating intelligent decisions from previously processed information, either from human or computational processes, is still dependent on support of the inferencing task.

In the event of a natural disaster, large amounts of information, both accurate and inaccurate, must be assessed by the ECC. The data being analyzed for the purposes of making decisions are reported to the ECC by *field agents*, or *sensors*. These field agents report such things as building damages, roadway damages, and civilian casualties. These sensors are disbursed throughout the disaster area and vary from trained professionals (police, EMT) and inexperienced civilians. The correct assessment of the information reported by these sensors is necessary to support the wide variety of authoritative decisions that must be made by the emergency command center (ECC).

For example, when a major natural or man-made disaster occurs, the ECC must make many assessments in order to make decisions, such as:

- Determining the immediate needs and priorities of the disaster victims;
- Determining the damage to critical facilities and civil infrastructure;
- Identifying stoppages, i.e. obstacles or interruptions to emergency operations or impediments to relief efforts;
- Monitoring public health; determining the resources available to respond to the disaster and identifying the gaps that need to be filled from outside resources.
- Monitoring public health facilities; determining the capacities and efficiencies of available hospitals and shelters for proper dispersion and assessment of casualties.
- Determining the location and availability of disaster-relief agents;
- Determining the damage to critical communication and transportation infrastructures necessary in order for disaster-relief agents to effectively assist in recovery.

The ECC's ability to properly fuse data plays a large part in the correct assessment of the disaster area situation. This in turn, is directly correlated to how well the ECC reacts to the situation defined by the natural or man-made disaster.

Before fusion can take place, the volumes of information arriving at the ECC from a series of disparate sources need to be modeled. This paper also focuses on the creation of the intelligent agent based reporting simulation that produces the ground truth and field reports necessary for data fusion to analyze. This reporting simulation is a crucial first step in the construction of a system wide simulation to produce a data fusion support tool.

Figure 1 illustrates how the ECC uses data fusion to produce a state estimate of the simulated environment. This situation estimate is based on reports received by field sensors. The fused data is used to make decisions regarding the direction of field resources.

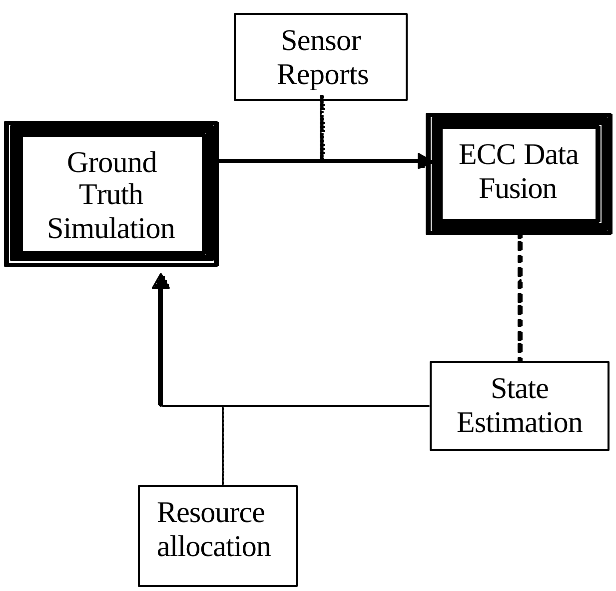

Figure 1: Ground truth simulation and sensor reporting with data fusion state estimation feedback

2 Scope of the Project

Several federates are collectively joined to produce the data fusion support tool. The focus of this paper is the intelligent agent sensor movement and report generation; however it is important that the reader have a clear understanding of the overall environment. The simulation has several core federates that each intake precise information, provide a processing function, and then output to the whole system. HAZUS, a custom Federal Emergency Management Agency (FEMA) application, is used to initiate the simulation and provide a preliminary picture of the post-disaster damages. The *ground truth* database maintains the real-time information regarding casualties and the level of damage in the disaster area. The *hospital module* simulates the fluctuations in the operating capacity of hospitals within the disaster area and updates the ground truth database. The *ambulance dispatching* federate uses a routing algorithm to efficiently distribute medical first-responders. These first three federates continually process and update their respective portions of ground truth throughout the lifespan of the simulation. As the ground truth is updated, the *sensor reporting* routinely takes this information from ground truth and creates a series of reports from a variety of sensor types to be sent to the data fusion federate. This reporting mechanism indirectly relays the information provided by the hospital module and the ambulance dispatching federate. The *data fusion* federate intakes the sensor generated reports and creates an estimate of the system. This system estimate is updated regularly as sensor reports continue to flow into the data fusion component. Figure 2 below gives a graphic overview of the entire system. Each of the core federates and some additional minor modules are displayed. In addition, the flow of information between each of the federates and the lesser modules is indicated by the direction of the arrows.

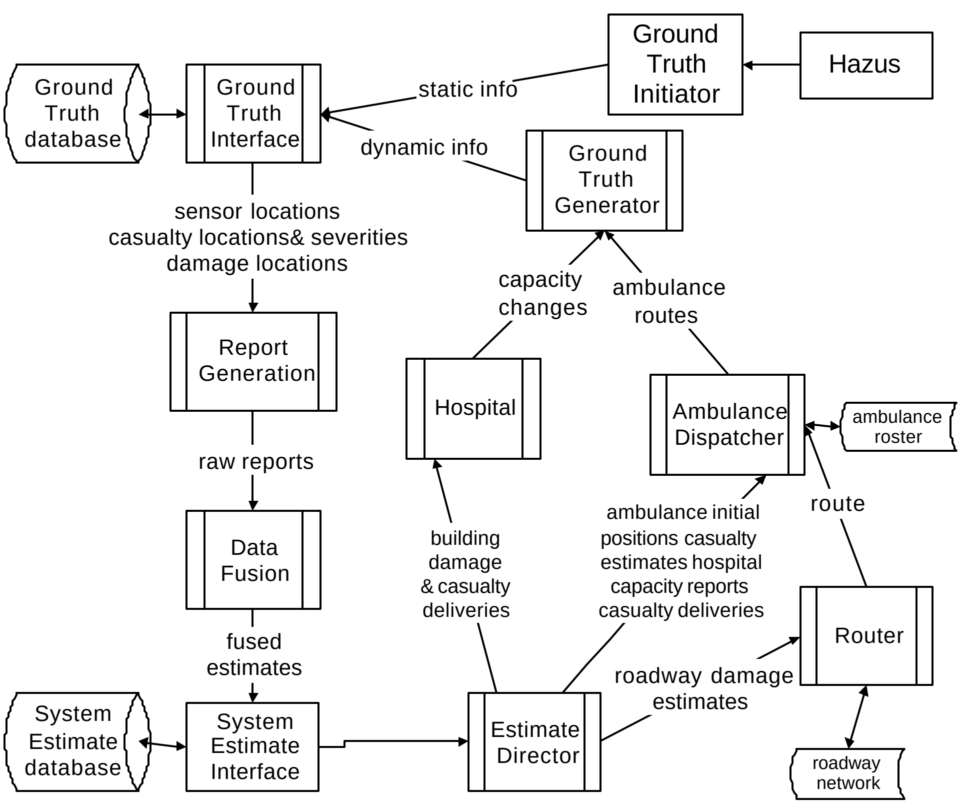

Figure 2 - Simulation Architecture for Creation of a Data Fusion Decision Support Tool for the Assessment of Manmade and Natural Disasters

As evidenced above, the sensor reporting function is crucial to the overall data fusion process. The remainder of this paper discusses the implementation of the intelligent agent reporting sensors.

3 Exemplar

The project exemplar used to illustrate the above system architecture is a simulation of the infamous Northridge Earthquake that occurred on the 17th of January, 1994. The Northridge quake occurred at approximately 4:30am PST, 20 miles Northwest of Los Angeles. The earthquake measured a magnitude of 6.7. Though the earthquake occurred early in the morning, significantly reducing the potential impact, 57 fatalities occurred and over 1,500 persons were "seriously injured" [2]. In addition, thousands of residents were left for days without electricity, gas, or water. Significant damage was made to the underlying infrastructure. Structural damage to buildings throughout the area and major roads added to the post-quake confusion. In this instance, as in that of future disasters, there is an obvious need for efficient and strategic decision making by the Emergency Command Center (ECC) in order to allocate a limited number of aid resources throughout the disaster region. Figure 3 depicts the Northridge, CA region where the 1994 earthquake took place. The major roads, hospitals and general layout of the region are all illustrated.

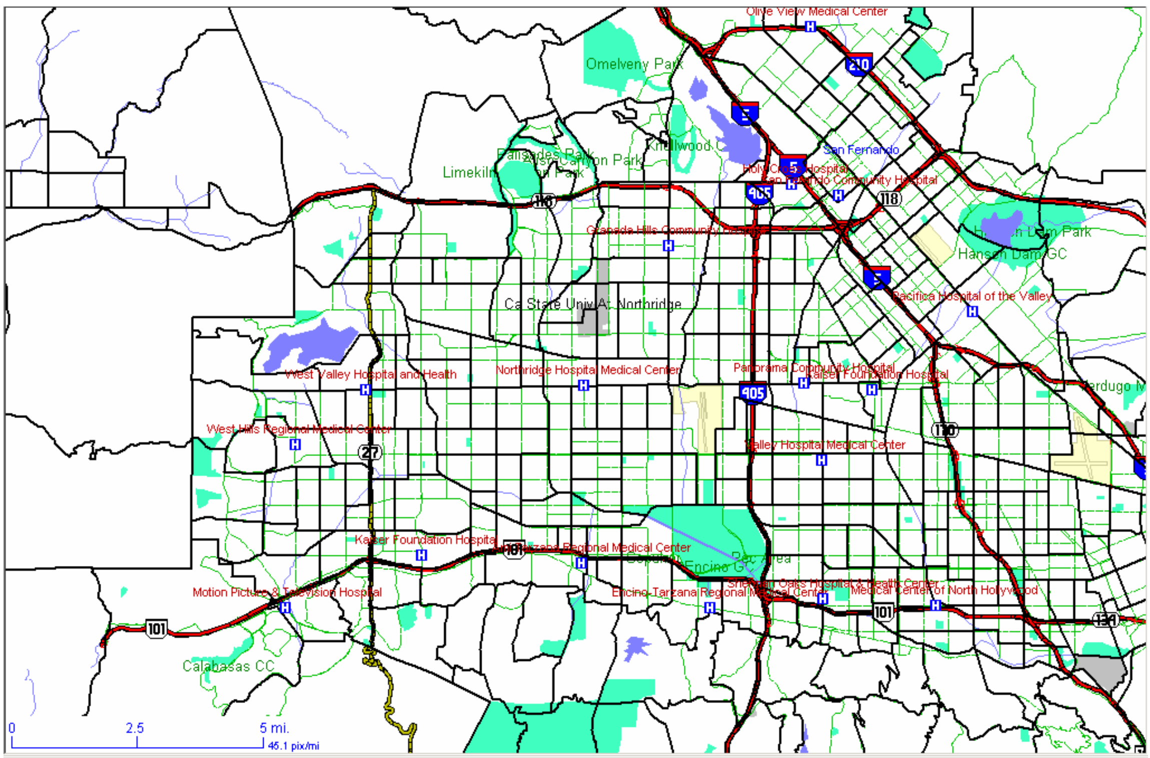

Figure 3 - Northridge Region

4 Software Requirements

The sensor reporting function utilizes the MapObjects geographical information system (GIS) development tool by ESRI. This application allows for the tracking of both spatial and characteristic information as well the graphic display of the sensor movements throughout the simulation.

5 Report Generation

The Report generator generates reports for essential facilities, casualties, and roadway - transportation systems. Reporting sensors including, police, civilians, ambulances and hospitals, issue the reports. Hospitals and civilians are treated as stationary sensors while police and ambulances are assumed to be mobile. These reports contain information regarding the location and the severity of damage and injuries resulting from the earthquake.

Several pre-defined report formats are simulated by the system in order to replicate the type of information that the Emergency Command Center (ECC) might expect to receive. Five basic report types are incorporated into the system. These include: essential and non-essential facilities, casualties, transportation systems, police and ambulance locations, and hospital status. Each of these report types are detailed below.

5.1 Essential and Non-Essential Facilities (Damages)

Essential and non-essential facility reports are issued by police sensors as they observe each facility. Essential facilities include hospitals, clinics, labs, blood banks, police, and fire stations. Essential facilities are noted in HAZUS prior to the deployment of the earthquake scenario. Non-essential facilities include residential, industrial and commercial buildings.

Facility damage is measured in three categories. *Structural damage* is damage to the facilities overall supportive structure. Acceleration-sensitive damage *is damage to ceilings, equipment that is an integral part of the facility such as mechanical and electrical equipment, piping and elevators.* Drift-sensitive damage *is damage to partitions, exterior walls, ornamentation and glass [3].*

Qualified staff working at each location makes reports for damaged essential facilities immediately. Reports made by hospitals and clinics include a functionality measure as well. The functionality measure is reported for each day; as a number of available beds remaining as a result of quake damage, and the number of beds expressed as a percentage of the total number of beds in the facility. Functionality of police and fire stations are also given as a percentage of full operational-functionality.

In the early stages of the simulation, only police report damage for essential facilities. These reports contain damage classifications for only the structural damage of the facility. In later iterations of the simulation, structural engineers will report damage for both non-essential and essential facilities. The reports made by structural engineers will contain highly accurate damage classifications for each of the three damage categories.

5.2 Casualties

In the current simulation, police, ambulances, and civilians report casualties. The police and ambulances acting as moving censors and report casualties as they either travel through an area of arrive at an assigned destination. They majority of casualty reports are sent by civilians who act as stationary sensors and call information into the Emergency Command Center (mirroring post-disaster 911 call volume). Casualty reports contain a specific patient ID, a location, and a classification of the severity of the injury. There are five classifications for the severity of the injury including: (1) basic medical aid – no hospitalization, (2) hospitalization – not life threatening, (3) life-threatening, (4) instantaneously killed or mortally injured, and (5) unknown.

These reports can be thought of as "observed casualty reports" in the sense that not all the casualties reported on were attended to by the sensor. For example, a police sensor may see a casualty while traveling to another location. The police sensor will report on that casualty but will not stop to attend to it, therefore omitting some information in the report such as casualty ID. The report will contain information such as casualty type, location, etc. (see report types outline).

5.3 Transportation Systems

Transportation system damage is reported for Roadways, Bridges, and Tunnels. In the current simulation, only the police and ambulances report damage to transportation systems. The damage to transportation systems is reported as slight, moderate, extensive, and complete. Reports for *roadway* damage contain the name of the road, the location, and the damage state. Reports for *bridge* or *tunnel* damage contain the name of the bridge or tunnel, the location, and the damage state.

5.4 Police and Ambulance Location

In the current simulation, police and ambulance sensor location reports will contain the location of the sensor and the availability of the ambulance (i.e. whether it can be dispatched or not). The information in these reports will be 100 % accurate. In later versions of the simulation this accuracy level can be adjusted to incorporate some amount of anticipated error.

5.5 Hospital Status

In the current simulation hospitals report their operating level and patient capacity. These reports are also currently 100 % accurate. Later versions of the simulation will have a function to perturb the data contained in these reports. Hospitals report operation level as a percent between 0 and 1. Where full hospital operation would be 1 (100%) and no hospital operation would be 0 (0%). The operation percentage should not be confused with capacity. Capacity is given as a function of beds. The hospital reports on the total number of beds it has and the total number of those beds that are occupied at the time the report is issued.

6 Sensor Reporting

As discussed above, police, civilians, ambulances, and hospitals, all issue reports that are then sent to data fusion. Each of these sensors not only reports on different information, but also operates in a slightly different manner. The operation of these sensors varies according to how we would expect the real-world entity to act during an actual disaster.

6.1 Police

Police agents are treated as intelligent agents who move from a starting location to a pre-defined destination. Prior to the start of the simulation, starting and target locations are assigned to each police vehicle. These starting locations can be either at a police dispatch area or anywhere in the region where police regularly patrol. The targeted end location can either be selected randomly or chosen to be a hot spot where excessive damage and casualties have been found. As the police sensor moves throughout the regions, they initiate their reporting on the status of their surroundings. A police sensor reports on essential facilities, casualties, and transportation systems. When a police sensor moves, it surveys the damages (casualties, facility structural damage, transportation systems) and reports. The police sensor senses reportable objects using a *distance to* method. If a reportable object (casualty, essential facility, etc…) is within a certain distance from the police sensor it is reported on. The distance that each police sensor uses can differ and is measured in meters from the point of the police sensor in all directions. For the current iteration of the simulation, the probability of detection = 1, meaning that the police sensor will report on all damages and causalities in its *distance to* buffer area. Initially this distance is set to 145 meters. Figure 4 below illustrates this reporting buffer area.

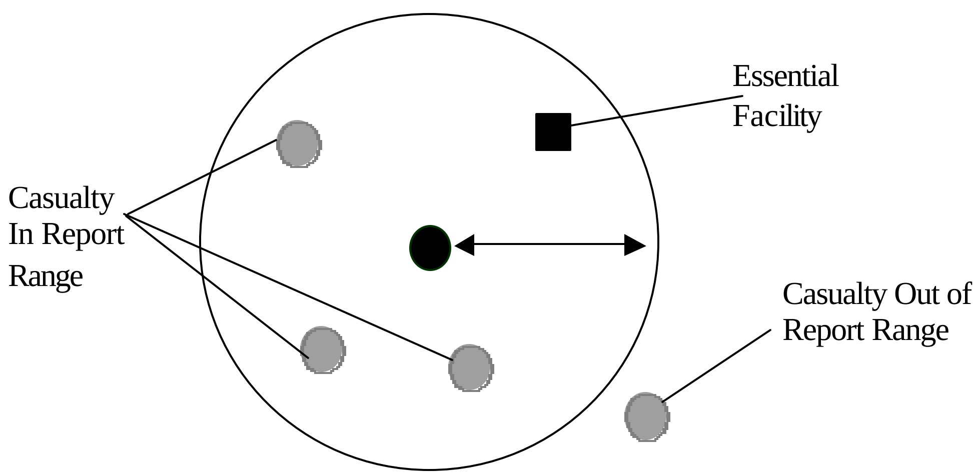

Figure 4: Moving Sensor Report Range

A single police sensor is not allowed report on the same damage or casualty twice. A second police sensor, however, can report on a damage previously reported by an initial police sensor. In addition to damage and casualty reports, police report on their location every five minutes. These reports contain the police's ID and their current location. These reports are meant to replicate the use of GPS devices in police squad cars that automatically send location information to the Emergency Command Center.

6.3 Ambulance

Ambulance sensors report on observed casualties, picked-up casualties and dropped-off casualties. In addition, ambulances produce reports when they become *idle* or when they become *stuck* due to roadway congestion or damage. Ambulance movement is directed by the dispatching module and a routing algorithm.

6.3.1 Observed Casualties

Observed casualty reports made by the ambulance sensor are similar to those made by the police sensor. These are reports made to resemble observation rather then examination and are assumed to be less accurate than pick-up and drop-off reports.

6.3.2 Pick-Up Report

Pick-up reports are made by the ambulance sensor when it arrives at its assigned location and picks up a specific casualty. These reports are assumed to be very accurate. Pick-up reports contain a field that indicates whether or not a casualty was picked up. In some cases the ambulance may arrive at an assigned location only to find that casualty is no longer there. In such cases the ambulance sends out a pickup report indicating that no casualty has been picked-up and that the ambulance is therefore ready for reassignment.

6.3.3 Drop-Off Report

Drop-Off reports are made by the ambulance when it arrives at a hospital and drops off a casualty. These reports contain the same information as the pick-up report and are primarily used to inform the emergency command center of the ambulances whereabouts and their availability for dispatch.

6.3.4 Location

Much like police sensors, ambulance sensors report on their location every five minutes. These reports include the ambulance ID, and the ambulance's location.

6.3.5 Stuck

The stuck report is made when the ambulance is unable to move de to roadway damage. This report contains the ambulance ID, the roadway ID (of the damaged road), the ambulance location and whether or not the ambulance is loaded. Loaded refers to whether or not the ambulance is carrying a patient (traveling to drop of the patient), or is without a casualty patient (traveling to pick up a patient).

6.3.6 Idle Awaiting Dispatch

The idle awaiting dispatch report is generated as soon as an ambulance becomes available for dispatch. Ambulances become available when the simulation commences, after each drop off, and when the ambulance reaches a pickup destination and the patient is no longer present in that location.

Figure 5 depicts ambulance states and the corresponding reports to be generated. Ambulance status is noted in italics and the reports are boxed adjacently. Loaded and unloaded indicate whether or not the ambulance contains a casualty.

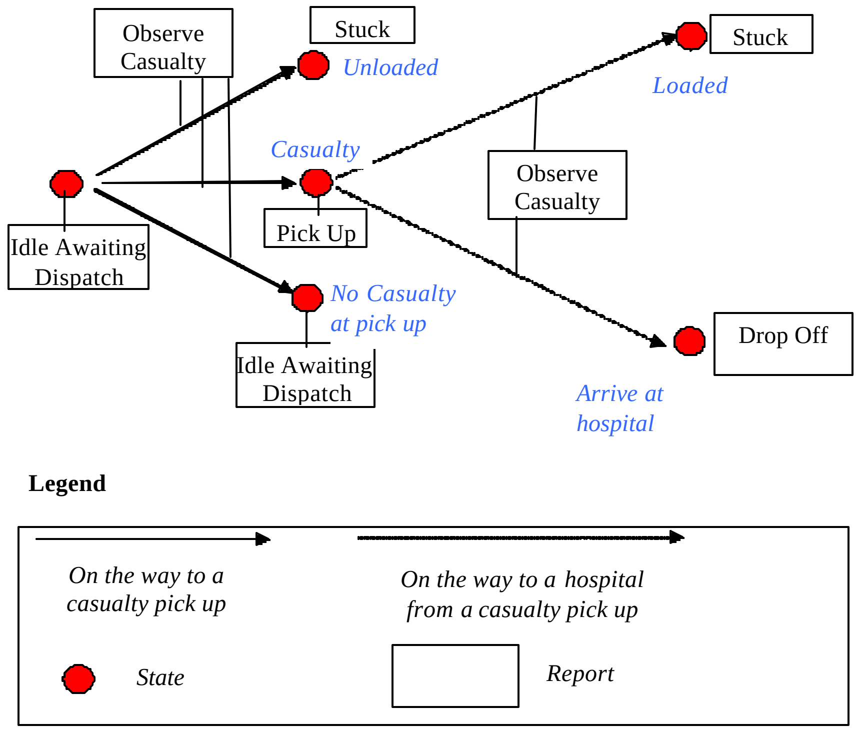

Figure 5: Ambulance Reporting Tree

6.4 Civilian

In the current simulation a civilian sensor is only enabled to report on casualties. These casualty reports are viewed as a random sampling and are issued in volume proportionate to population in the area surrounding the casualty (census tract) as well as a time decay function. Each civilian reports on a random casualty. Because the reporting is randomized, more than one civilian in each census tract may report on the same casualty. The format of the civilian-observed casualty report is similar to the police and ambulance reports. Because non-casualty civilians are not tracked in the ground truth, a generic Report ID value (0) is used for all reporting civilians. Additionally, civilian reporters are assumed to be relatively stationary so no mechanism for movement has been assigned to them.

At present, immediately after the earthquake (time less than 10 minutes) only 0.005 % of the population in each census tract produces a casualty report. The casualty reported on is chosen at random from the pool of available casualties in the tract. Each civilian reporter reports on only one casualty at a time. In the early aftermath of the earthquake (between 10 and 30 minutes), 0.0025% of the population in each census tract files a casualty report. Again the precise casualty is chosen at random with each civilian reporter only reporting on a single casualty. After 30 minutes, no civilian reports are made. Additionally, a rule is instantiated to disallow more than three reports on any single casualty. This rule is aimed at reducing the likelihood of over-reporting for tracts that have a large population of potential reporters and a low number of casualties.

6.5 Hospital

The hospital sensor reports on the operation level and the capacity of the particular hospital that is filing the report. Each hospital reports separately and independently. The reports are made consistently every thirty minutes.

7 Reporting Caveats – Incorporating Error

As discussed above, all reports are created using a standard format defined by the system. In order to model some of the chaos and error that exists in the real-world, the standard formatted reports should have some missing or inaccurate information incorporated. Building in this error can be achieved by randomly removing information from full reports and confusing some of the existing accurate information. Additionally, the amount of expected error varies based on the sensor type. Civilians for example, have little expert knowledge therefore we should expect their reports to be the least accurate. Police and ambulance sensors should have significantly more expertise in casualty damage estimation and therefore, we would expect only a limited amount of error in their situational assessments. Structural engineers that will be incorporated into later versions of the simulation should have only very minor if any inaccuracies in their reports on building and transportation system damage. The reports, made by the field sensors regarding the levels of severity of injuries and damages, will vary in error according to the sensor type (police, ambulance, civilian) and the object being reported on (facility, casualty, transportation system). Confusion matrices will be used to reflect the level of confusion of each field sensor. Each field sensor type will have a distinct confusion matrix for each targeted report object. These matrices will represent the probabilities of the sensors correctly assessing the true state of the target object.

In addition to the severity of injuries for casualties and the level of damage for structures, the locations of the reported entities are also confused. This error can be viewed as either the result of incorrect reporting or as the distortion of information as it is relayed from the original reporter through the information network. The location in each report is confused using a Gaussian distribution around the actual location. This method confuses both the x and y locations in the report. In the current simulation, the standard distribution for the police reporters is calibrated to 10 meters while the standard deviation for the civilian reporters is set to 20 meters.

Conclusions

The current agent-based reporting system allows for the creation of a myriad of reports to be generated for input to data fusion module. This reporting mechanism has been created to approximate the actual type and amount of information that an emergency command center might expect to receive in a post-disaster environment. The reporting system incorporates several first responder agents as well as civilian and hospital originated information. These reports have further been distorted to capture a realistic quantity of error incurred by either lack of expertise or gaps in communication. These reports are then made available to the system wide simulation for data fusion efforts. The reporting efforts and data fusion continue in parallel for the life of the simulation. This reporting simulation is crucial to the overall post-disaster simulation as fusion can not be attempted until lifelike input is available.

Though this paper has focused on a post-earthquake scenario, the intelligent agent based reporting simulation and data fusion are applicable to a wide array of manmade and natural disaster scenarios. This type of reporting sensor simulation with data fusion is a first step toward the objective of creating fast, accurate, understandable and trustworthy estimates necessary to support a variety of decision-making requirements in post-disaster settings.

References

[2] Aurelius, Earl, "The January 17, 1994 Northridge, CA Earthquake," [Online document] March 1994, [2003 May 8], Available at HTTP: http://www.eqe.com/publications/northridge/northridge.html

[3] Scandale, Jim. "A Review of HAZUS99 Output," May 2002.

Back to the Top