

Event-driven, Role-based Mobility in Disaster Recovery Networks*

Samuel C. Nelson, Albert F. Harris III, and Robin Kravets
Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, IL, USA
snelso20@cs.uiuc.edu, aharris@cs.uiuc.edu, rhk@cs.uiuc.edu

ABSTRACT

One of the most important tools in understanding the complex characteristics of disaster recovery networks is simulation. While many mobility models exist for simulating ad hoc networks, they do not realistically capture the behavior of objects in disaster scenarios. We propose a high level event- & role-based mobility paradigm in which objects' movement patterns are caused by environmental events. The introduction of roles allows different objects to uniquely and realistically react to events. For instance some roles, such as civilian, may flee from events, whereas other roles, such as police, may be attracted to events. Furthermore, to incorporate reaction from multiple events in a realistic fashion, we propose a low-level gravity-based mobility model in which events apply forces to objects. Simulation results show that our disaster mobility paradigm coupled with our gravitational mobility model creates a network topology that differs from the popular Random Walk mobility model. This new disaster mobility model opens up the door for more realistic simulation of communication and routing protocols for disaster recovery networks.

Categories and Subject Descriptors: C.2.1 [Network Architecture and Design]: Wireless communication

General Terms: Design

Keywords: Disaster networks, mobility modeling, simulations

1. INTRODUCTION

Much of the recent interest in delay tolerant networks has come from the need to support communication for organizations like police and fire departments, as well as other first responders. The behavior of most of these organizations is driven by the need to respond to events and participate in those events based on the particular role of the organization. Since this type of behavior is very specific to emergency and

*This work was partly funded by Boeing and NSF CNS 0347468.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHANTS'07, September 14, 2007, Montréal, Québec, Canada.
Copyright 2007 ACM 978-1-59593-737-7/07/0009 ...\$5.00.

disaster response scenarios, understanding communication patterns in such networks is critical to understanding how to improve the current state of emergency response. However, current mobility models for wireless networks do not capture the complexity of either the behavior of the different components of such networks or the specifics of the expected communication patterns. Realistic mobility and communication models can enable more effective evaluation via simulation and eventually lead to more effective solutions.

One of the biggest challenges faced by communication in disaster recovery networks comes from the high expectation of network partitions and dynamic object behavior. The high potential for object failure further complicates matters. Intuition says that, as objects react differently to different events, the variance in density of the communication graph will change and cause strain on current routing protocols. Due to this unique behavior, disaster recovery networks present challenging environments.

One of the fundamental aspects of simulating disaster recovery networks is realistically modeling the movement patterns of the mobile objects. Modeling mobility enables testing the effectiveness of current routing protocols as well as provides insight into how routing protocols can be improved. In a disaster environment, it presents unique challenges in that environmental events and roles directly affect a node's movement patterns. Intuitively, events act as stimuli for mobile nodes in the network, causing them to react in ways according to the predefined roles they take on. Many roles in disaster networks must react to multiple events by fleeing or approaching in a realistic fashion.

Many ad hoc network mobility models have been developed and analyzed [4, 1, 5, 7, 9, 11]. Models based on random movement are particularly popular and heavily studied [2, 3, 13, 6]. These models, while adequate to study environments for which they were designed, do not allow for groups of objects to react differently to environmental events. Therefore, a higher-level, more general mobility model is needed to incorporate the different roles objects play in disaster scenarios.

In this paper, we present an event- & role-based mobility paradigm that effectively characterizes the movement patterns of objects in a disaster scenario. Different sets of these patterns are embedded into different object roles. Attaching actions to roles and not directly to objects has the advantage that movement patterns are organized and objects can quickly shift from one pattern to another by assuming different roles. To the best of our knowledge, this is the first disaster mobility paradigm that is reactive, in a role-based

fashion, to environmental events and their associated parameters.

The main contribution of our research is the classification of a generic event- & role-based mobility paradigm that completely defines movement patterns given a series of environmental events for a set of characteristic roles. Additionally, we present a low-level physics-based gravitational mobility model that “plugs in” to our event-driven, role-based paradigm allowing objects to react to the presence of numerous disaster events based on the particular role of the node. This allows objects to flee from or approach multiple, unrelated events. To evaluate the effect of our comprehensive model on communication patterns in disaster recovery networks, we discuss a new set of relevant metrics that help characterize the changes in topology as disaster events unfold. Finally, we have developed a set of tools to realistically construct a mobility scenario and trace files of our disaster mobility model for the ns2 network simulator [10].

The remainder of this paper is organized as follows. Section 2 discusses the properties of objects in disaster recovery scenarios that affect their mobility patterns. Section 3 presents the details of our event- and role-based mobility paradigm. Section 4 presents a gravity-based low-level mobility model that can be used with our high-level paradigm. Section 5 presents the methodology we use to study the properties of our model, including a presentation of the metrics used to evaluate the resulting networks. Section 6 presents the results of our simulation study. Finally, Section 7 presents some conclusions and future directions.

2. MODELING OBJECT BEHAVIOR

Modeling the movement behavior of mobile objects has been heavily studied. Due to its simplicity and effectiveness, the Random Walk [6] model is widely used to model such object behavior. In this model, a node randomly chooses a direction in $[0, 2\pi)$ along with a speed, and moves according to those choices for a set amount of time. After this time has expired, the node repeats the process. A recent model by Jardosh *et al.* [7] takes polygon-shaped objects into account by using Voronoi diagrams to build walks. While many of these models are adequate for their particular environments, all objects generally act in the same way and, therefore, do not give the flexibility required for modeling disaster scenarios. This is because real people and vehicles take on roles, allowing them to react to events in a distinct fashion.

Mobility in disaster recovery scenarios is fundamentally driven by environmental events. These events act as stimuli towards objects and directly cause them to change their movement patterns. While some current models, including [7], could be extended to react to external stimuli, truly capturing the complex interactions between more than one environmental event requires building new mobility models with event-driven actions as the primitive concept. Furthermore, while all objects react to relevant events, different classes of objects react in different ways. In other words, object behavior changes over time and is not uniform across all nodes. Objects also must react to multiple events in a realistic and smooth fashion. There is currently no adequate mobility model that takes into account these observations.

To illustrate this idea with an example, consider an apartment fire in a populated neighborhood. There are, intuitively, at least three different classes of behavior that objects can assume: (1) fleeing from the event, as is the case

with civilians, (2) approaching the event with the intent of staying, as is the case with police and firefighters, and (3) oscillating from event to a predetermined location, as is the case with ambulances. These high-level behavior classes, which we refer to as roles, help give general but clear mobility patterns, which are realistic and relatively easy to simulate.

Roles, however, need not be limited to specifying movement patterns during a disaster scenario. It is easy to extend the concept of a role to cause an object to react differently during different stages of an event. For example, cleanup crews may want to react to disaster events by approaching them long after the event has occurred, whereas police may instead want to approach the event immediately.

In response to the need for unique behavior modeling in disaster scenarios, we have developed a high-level, role-based, event-driven mobility paradigm in which different roles take on different mobility patterns in reaction to specific events. Our paradigm is high level in the sense that an object may take on multiple mobility patterns over the course of some time period. For instance, a civilian may first be modeled by Random Walk. After some time, an event, such as a fire, may trigger the civilian to change its model to one of fleeing from the fire. By modeling mobility as a series of event-driven, role-based actions, we can properly select which specific mobility model to use for a given object and situation. While the specific rules for reaction should be based on observational studies, our paradigm is sufficiently general enough to allow future, accurate movement patterns to be used.

Our paradigm allows for different objects to react to events in unique ways by attaching a mobility pattern for each possible (event, role) combination. These lower level mobility patterns can be any of the previously developed mobility models, including Random Walk. To further illustrate the behavior of certain roles in a disaster scenario, we have developed a low level physics-based gravitational model that allows objects to flee or approach disaster events in an intuitively realistic fashion. It is important to note that our high level, event- & role-based paradigm is not tied to the gravitational model in any way, and can support any of the previously defined low level mobility models.

Our gravitational model captures the effect-distance relationship between events and objects, allowing events to act on objects via forces. Objects closer to events are affected more than objects farther away. Additionally, each event has an *event horizon*, defined as the maximum distance at which an event affects objects, which allows events to have a defined radius of effect. Finally, there is a *communication threshold*, which defines the time until emergency vehicles are notified of an event. Until this time, emergency vehicles outside the event horizon do not respond to the event; however, after the threshold time has passed, those emergency vehicles begin to converge on the event. It is possible that multiple events could take place in a single scenario. Additionally, these events may or may not be simultaneous. One of the major benefits of the gravitational model is that it easily captures the interactions between multiple events.

The classification of behavior into roles can also play a part in establishing realistic communication patterns. For instance, civilians are most likely in contact only with police and other civilians, police are in contact with all roles, firefighters are in contact with police and other firefighters,

and ambulances are in contact with police. These communication patterns, along with the mobility, can simplistically model an entire disaster scenario. For this paper, however, we concentrate solely on modeling mobility,

3. DISASTER MOBILITY PARADIGM

In this section, we formally describe our high-level mobility paradigm that incorporates external, environmental events along with role-based reaction. By partitioning objects into roles, which define their reactions to events, our paradigm can realistically obtain a set of (role, event, action) triples that define the overall movement patterns of objects in disaster recovery scenarios. A single triple can be read as follows: “Role r reacts to event e by taking action a ”. Then, by instantiating the triples with the characteristics for different agents operating in a disaster scenario, a mobility pattern can be generated for the scene.

Three entities, and their relationships to one another, help define our high level mobility paradigm: objects, roles, and events. *Objects* are nodes in the system that provide movement and communication. Each object assumes a *role*, or set of roles, that indicate what movement pattern the object should assume in response to external stimuli. The specific areas of interest that provide external stimuli to objects are referred to as *events*. The event-based response a role dictates to an object is an *action*, which is generally a low-level mobility model, such as Random Walk or the gravitational model presented in Section 4.1.

We now elaborate on these three entities and their relationships to one another.

Objects are the critical components of the scenario, including people, buildings, and vehicles, and are defined by the following parameters:

- Location: The (x,y,z) location of the object.
- Role: The role, or set of roles, associated with the object.
- Velocity: The current velocity of the object (in vector form).

Roles dictate how objects react to events. We define four main categories of roles, although it is possible to define any number of them. First, the *repelling* category causes objects to be repelled from events. The low-level mobility models that support this role should allow objects to move away from events in a realistic and easy-to-use fashion. The most common use of this role is to model normal civilians in a disaster scenario. An attribute of this role is *curiosity*, which dictates how likely it is for an object to stop at the event horizon, simulating curious on-lookers. Second, the *attraction* role causes objects to converge on events. Low-level mobility models that support this role should cause objects to quickly approach an event or events. Common uses of this role are to model police and firefighters. Third, the *oscillating* role models objects that first approach an event and then, upon reaching the event, travel immediately to a pre-defined location. This movement pattern is then repeated continuously. Low-level mobility models that support this role should allow this action to be as realistic as possible. One use of this role is to model an emergency response system in which ambulances oscillate between the event and a hospital. Finally, the *immobile* role models any object

that remains stationary for the duration of the simulation or until the object takes on a new role. This role can be supported by the lack of a low-level mobility model, since it does not perform any movement. This role is useful to model both naturally static objects, such as hospitals, as well as event-caused immobility.

We anticipate that the default action dictated by many roles is Random Walk, since it simply models motion when movement patterns are unknown or seem random. It is important to note, however, that any mobility pattern, such as one that accounts for navigating around buildings or objects, can be used.

Events act as the stimuli for mobility changes in the scenario. In a real disaster scenario, an object’s proximity to an event is a major factor in how it reacts. Therefore, it is important to clearly mark distinct areas of an event. Our paradigm captures this type of behavior by defining the *Disaster Radius*, the *Event Horizon*, and the *Relevant Radius*. Reaction to an event is also dependent on time, which is modeled by the *Start Time* and the *Radio Contact Time*. The following defines the full set of parameters for an event.

- Location: The (x,y,z) location of the event.
- Start Time: Time when the event occurs.
- End Time: Time when the event ends.
- Radio Contact Time: Time when radio contact outside of the event horizon occurs.
- Disaster Radius: Area inside which all objects become immobile.
- Event Horizon: Area inside which all objects react to the event, even before radio contact occurs.
- Relevant Radius: Area inside which objects, based on their role, react to the event after radio contact occurs, assuming they are not already reacting. Some roles, such as the civilian role, may not react when inside this radius but outside the event horizon.
- Intensity: A numeric representation of the event’s current intensity.

4. A DISASTER MOBILITY MODEL

With the high-level disaster mobility paradigm formalized, we now describe a disaster mobility model that we believe intuitively models simple disaster scenarios. For implementation purposes, we have made some simplifying assumptions. First, we assume events are stationary, have a constant intensity, and, after their start, persist for the remainder of the simulation. Furthermore, objects assume a single initial role and do not change it for the duration of the simulation, unless changing to the immobile role. Finally, all roles except the Immobile role default to the Random Walk action. At the end of this section, we discuss how to capture events that are mobile and change intensity, as well as events that are not best represented by a single point.

4.1 Gravitational Model

Intuitively, many roles react to disaster events by either fleeing from or approaching them. To model these actions in the presence of multiple events, we use a physics-based

gravitational model to define the “Flee” and “Approach” actions. Gravitation has been used to model group mobility dynamics, particularly in [14], but not event-based mobility. We designed this model based on the observation that objects, in general, either gravitate towards or away from disaster events.

Physics states that the gravitational force between two objects with masses m_1 and m_2 at a distance d from each other is:

$$F = \frac{G \cdot m_1 \cdot m_2}{d^2}$$

where G is the gravitational constant. The total resulting vector force on an object in the vicinity of multiple objects is calculated by the vector sum of all forces on the object. The resulting force directly affects the object’s acceleration.

We borrow the concepts of gravity and force from physics to model the *flee* and *approach* actions. By letting m_2 be the “mass” of a given event, and assuming m_1 is negligible, the force on an object by that event can be described as $F = I/d^2$, where I , the intensity, encompasses G and m_2 . Mobile objects can then be repelled or attracted to events (or multiple events) by assigning a particular intensity to every event.

To calculate the motion trajectories of objects as a result of multiple forces, we sway from physics slightly to allow for a more intuitively realistic movement pattern. Physics states that forces directly affect an object’s acceleration. However, humans are more concerned with maintaining a particular velocity at a given time than maintaining a particular acceleration. We generally do not maintain a constant acceleration, but rather accelerate quickly to a desired velocity and then hold an acceleration of zero. Therefore, it is intuitively more correct to say that humans will adjust their *speed*, not their *acceleration*, according to how far they are from a disaster event (of course, they will adjust their acceleration to obtain that speed, but only for a short period of time). Therefore, in our gravitational model, forces act directly on velocity, not acceleration, to account for this phenomenon. The benefit of taking a gravitational-based approach to model mobility is that it allows the reaction of objects to be intuitive and easy to compute computed for multiple, unsynchronized, dynamic events.

4.2 Disaster Model

Let M be the set of (role, event, action) triples that define our mobility model. M is populated by adding triples to cover all components or the desired scenario. All mobile nodes initially start using the Random Walk Model to either walk or drive, with speeds appropriately bounded. Formally, there are a set of initial (role, event, action) triples for each role as follows:

$$\{(r, \text{No Event}, \text{Random Walk}) : r \in R\},$$

where R is the set of all possible roles. The “No Event” event is simply the default event every role assumes when there are no relevant events in the network.

If a disaster event occurs, two areas are immediately formed. The first area is ground-zero, as defined by the disaster area parameter, within which objects are immediately immobilized. We model this by simply immobilizing all nodes within a set radius of the disaster event. We formally model this by the inclusion of the following triples into M :

$$\{(r, \text{DE1 - At Ground-Zero}, \text{Switch to Immobible}) : r \in R\}$$

The “DE1 - At Ground-Zero” event is a disaster event (with the label “DE1” representing disaster event #1) that has occurred when the object was within the disaster radius of the event. Once an object is immobile, it stays immobile for the remainder of the simulation. To accomplish this, the action “Switch to Immobible” instructs the object to immediately switch roles to the “Immobible” role. This role is defined in M as follows:

$$(\text{Immobible}, \text{No Event}, \text{Stay Still})$$

It is important to note that this should be the only entry for the immobile role, since it should always default to staying still. Static objects, such as a hospital, are also assigned the immobile role.

The second area is defined by the *event horizon*. All objects within the event horizon react to the event by either gravitating towards it or fleeing from it, at a speed dependent on the object’s proximity to the event. The inclusion of a set of triples into M formally models this phenomenon. For instance, the following triples define the area within the event horizon for a simple disaster scenario:

$$(\text{Civilian}, \text{DE1 - In Event Horizon}, \text{Flee})$$

$$(\text{Police}, \text{DE1 - In Event Horizon}, \text{Approach})$$

$$(\text{Firefighter}, \text{DE1 - In Event Horizon}, \text{Approach})$$

$$(\text{Ambulance}, \text{DE1 - In Event Horizon}, \text{Oscillate})$$

The event “DE1 - In Event Horizon” event refers to the situation that the object is within the event horizon radius of a disaster event. Notice that while all the events beginning with “DE1” are technically the same event, to incorporate proximity into the action taken by a role, we break the event into multiple areas (or regions), in which roles reacting to the same event may respond differently based on which area of the event they are in.

After the radio contact time of the event expires, the relevant radius of the event is formed. Roles with radio contact in this region, but outside the event horizon, should react to the event. Continuing from the previous example, the following tuples in M formally define this area:

$$(\text{Police}, \text{DE1 - In Relevant Radius}, \text{Approach})$$

$$(\text{Firefighter}, \text{DE1 - In Relevant Radius}, \text{Approach})$$

$$(\text{Ambulance}, \text{DE1 - In Relevant Radius}, \text{Oscillate})$$

An object may have multiple applicable triples at any given instance. This would occur, for example, if a civilian were in the radii of two different events. Our gravitational model easily accommodates scenarios of this type.

It is possible to extend our model to account for events of different shapes and sizes, as well as mobile events. Currently, events provide forces from a central point within the event, and have different radii that allow for a circular (or spherical) shape. Elongated event shapes, such as floods, can be simulated by placing multiple events close to each other at varying intensities. Furthermore, there is nothing prohibiting the changing of intensity or location of an event, as forces can be quickly recomputed at every object’s location based only on current information. The natural memoryless computation of forces allows for highly dynamic events.

5. ANALYZING MOBILITY MODELS

The benefits of an effective mobility model come from its ability to capture and expose the characteristics of the network and the behavior of nodes in the network. This information can then be used by network designers to understand how to design protocols that are suitable for the specific scenarios. In this section, we first discuss the metrics necessary to describe network behavior in disaster scenarios. Although most evaluations focus on simple network characteristics, such as node density and path length, the unique behavior of nodes in a disaster scenario results in more interesting network conditions that require us to look at more complex parameters, such as average node density and partitioning. We then present a set of tools that we implemented to generate ns2 mobility scenario files. In the next section, we present our evaluation of a number of these metrics using our tools.

5.1 Metrics

When discussing mobility models, it is important to understand how a model affects different topological network metrics. Two standard metrics are average node density and average path length. Average node density, as defined by the average number of neighbors per node, can be used to help characterize the potential connectivity of a network, since a network with low density will likely be partitioned. Average path length, as defined by the average number of hops from source to destination, captures the distance between sources and destinations. However, due to the highly dynamic nature of networks under disaster scenarios, it is important to not only consider these metrics but also those that show how the structure of the network progresses over time. For an event-driven, role-based mobility model, the following network parameters highlight how the network is changing over time.

- **Partitioning over Time:** The average path length metric is meaningless when the graph is partitioned, which is likely in many disaster scenarios. Therefore, tracking whether or not the graph is partitioned is critical to understanding the flux in network topology.
- **Clustering Coefficient over Time:** The clustering coefficient of a particular node is the number of that node's neighbors that are connected to each other divided by the total possible links between them [15]. The higher the value of this metric, the more clustered the graph is.
- **Average Node Density over Time:** Although node density is an important metric, in an event-driven mobility model, it is important to capture how density changes in *reaction* to different events, which indicates churn.
- **Maximum Node Density over Time:** Maximum node density gives insight into the potential cluster sizes, which can provide insight into potential bottlenecks. Although it would be useful to track actual cluster sizes, maximum node density provides a much cheaper, though quite effective, heuristic.
- **Variance of Node Density over Time:** Since some parts of the network may be more stable than others, the variance in node density gives insight into the amount of variance in cluster sizes as a result of different events.

5.2 Tools

To evaluate the impact of our model on the metrics described above, we have implemented two tools for the ns2 simulator. The first tool is a parameters file generator that creates a properly formatted parameters file appropriately choosing random values when necessary. This tool prompts the user for the following input: size of the network (in terms of meters squared), number of civilians, number of ambulances, and number of police. Since both the police and firefighter roles are similar, we have chosen to omit firefighters and simply add more police to simulate firefighters. However, it is quite simple to include firefighters, or other responders, and give them appropriate behaviors. The output parameters file contains the following information:

- Grid size and simulation runtime
- Randomized coordinates for all objects and events, and coordinates for four hospitals located at the corners of the grid
- Minimum and maximum speeds for objects
- Percentage of curious civilians
- Random Walk parameters
- Randomized trigger times and radio contact times for four events
- Randomized intensities for events, which determine radii for event horizons and damage zones

The specific parameters generated for the experiments in this paper are detailed in Section 6.1. For any given input, this tool produces unique output since many parameters are randomly chosen. Usage for this tool is as follows:

Usage: paramGen > paramFile.

The second tool is our main event-driven simulator. This tool accepts as input the parameters file generated by the first tool and runs a complete simulation with knowledge of Random Walk and our physics-based gravitational model. It is important to note, however, that any mobility model can be plugged into the tool in place of Random Walk and/or the gravitational model. The output of this tool is an ns2-compatible mobility trace file that gives the current velocity and destination of every object at every second in the simulation. All of the mobility model logic is performed in this tool. For any given input, the output of this tool again produces unique output, since Random Walks performed by objects not reacting to events may differ from one simulation to the next. Usage for this tool is as follows:

Usage: disasterSim [-d] < paramFile > nsMobilityTrace.

The *-d* flag, when passed, displays each individual step of the simulation via "ASCII art" to the console. This is printed to standard error, so it is not written to the *nsMobilityTrace* file.

Both of these tools are open source and will be available for download at our website.




Figure 1: Network snapshots at time 0, 200, 400, and 1400 seconds

6. SIMULATION RESULTS & ANALYSIS

To analyze the difference in network topology changes generated by our disaster recovery mobility model, we generated numerous topologies and ran simulations with them using ns2. Using the same initial setups, in terms of node placement and numbers, we ran the simulations using the Random Walk model for comparison. In this section, we present results from 10 different sets of simulations, each with its own group of both deterministically and randomly chosen parameters. Each of the metrics presented in Section 5.1 are evaluated for each of the resulting sets of trace files. Sufficient numbers of experiments were run to minimize the 95% confidence interval.

For this paper, we were only interested in the mobility patterns of objects in a disaster scenario and the topological effects the patterns have on the network graph. Therefore, we did not simulate communication between nodes. However, we did specify the communication range to be 150 meters to obtain information about links in the network.

6.1 Simulation Parameters

The simulation time runs from 0 to 1500 seconds with a grid size of 1000 m^2 . The communication range of every node is 150 m to simulate an urban environment. There are 75 civilians, 10 ambulances, and 15 police each randomly located. A total of 90% of civilians, randomly chosen, are considered curious, meaning they stop at event horizons to look at the event. Furthermore, we have chosen a minimum and maximum speed of 1 m/s and 4 m/s respectively for civilians, and 17 m/s and 20 m/s for ambulances and police. All Random Walks are done for 30 seconds.

Four events are randomly placed on the grid. The first event occurs randomly between 100 and 200 seconds, the second between 125 and 225 seconds, the third between 150 and 250 seconds, and the fourth between 175 and 275 seconds. Radio contact for a specific event occurs randomly between 40 to 80 seconds after the event has occurred. The intensity of events are randomly chosen between $10000 \text{ m}^3/\text{s}$ and $20000 \text{ m}^3/\text{s}$. The event horizon for each event is 2% of the intensity and the damage radius is 0.1% of the intensity.

We choose high intensity events to easily illuminate the differences between the topology of our model versus the topology of the Random Walk model. We also choose to include hospitals as stationary objects, since they will most likely participate in communication with other objects (particularly ambulances and police). Therefore, there are a total of 104 objects in the system, 4 being immobile from the beginning. Furthermore, we assumed that radio contact for all events reached ambulance and police regardless of where

they were. In other words, all responders are in either CB or cellular radio contact at all times, meaning that the relevant radius for each event is set large enough to encompass the entire grid.

The model we present contains a large parameter space. This is due to not knowing how the different parameters affect communication and routing in disaster events. As future work, we plan to explore the effects of these parameters and simplify the mobility model based on those findings.

6.2 Snapshot of Topology Change

Figure 1 shows a series of four snapshots during one simulation run of our disaster mobility model. The first box shows the state of the network at the start of the simulation. Object location at this point is random, except for the 4 hospitals located at each corner of the grid. The second box shows the state of the network 200 seconds into the simulation. At this time, some events have triggered but radio contact for many has not occurred. Only objects within the event horizon have reacted at this time. The third box shows the state of the network 400 seconds into the simulation. By this time, all events have been triggered and radio-contact has been made. All emergency response objects (police and ambulances) and civilians within the event horizon have reacted to the event. It is now possible to see some of the different roles active in the system, simply by visually observing their locations. The fourth box shows the state of the network 1400 seconds into the simulation. By this time, most metrics have come close to convergence and mobility is noticeable only by ambulances and civilians who have not approached the event horizon. Civilians who are not curious and have left the event horizon are mobile again.

A clear topological difference between the disaster mobility model and Random Walk is primarily due to the clustering of objects around the event horizon. This separates the graph into three primary areas: (1) areas inside event horizons, (2) areas at or near event horizons, (3) areas outside of event horizons. The first area is very sparse since all civilians are able to move leave the scene. Almost all of the concentration is in the “damage radius”, since emergency response workers immediately move towards that zone. The only interaction between that zone and the event horizon are the oscillators when they pass through the event horizon. The second area is very dense since all of the civilians inside the event horizon gravitate towards its edge and civilians that happen to stumble into the event horizon stay there. The third area contains objects who are performing random walks and have not been notified of the event. As the simulation continuously runs, the third area should slowly lose objects to the




Figure 2: Average node density




Figure 3: Maximum node density




Figure 4: Variance of node density

second area since they randomly hit the event horizon.

This series of snapshots clearly shows the location of events and the formation of crowds of people around the event horizon. It also illustrates the behavior of ambulances going to and from events and hospitals. We would expect very similar results in a real disaster scenario, further confirming our implementation.

6.3 Metric Evaluation

Figures 2 through 6 clearly show the event-driven response of the metrics around the time of the events. Between 0 and 100 seconds, the data sets are similar for all metrics, as expected. Between 100 and 355 seconds, during the triggering of events and radio contact time, a clear divergence between the disaster mobility model and Random Walk model is readily seen as the topology of the event-driven simulation starts to take form.

Figure 2 shows the average node density of the network as time progresses. The average node density in the disaster mobility model increases in response to events. This is due to the gathering of nodes around the event horizon, forcing them into a smaller area than before. An interesting observation is that the size and frequency of the oscillations in average node density become both smaller and less frequent in the disaster mobility model as time progresses. This is due to the topological convergence of the disaster mobility model that does not occur in Random Walk. The difference in average node density between the disaster mobility model and Random Walk is important because it gives overall information as to how many neighbors a node can expect to have at a given time and so provides hints about network connectivity.

Figure 3 shows the maximum node density between the two data sets as time progresses. The maximum node density quickly increases in response to the events. There are two highly-clustered areas for each event in the system, the area inside of the damage radius and the event horizon. The jump in maximum node density is due to the quick response by the police to the event, increasing the node density of nodes at the damage radius. After this, the maximum node density remains relatively constant from around 500 to 900 seconds, as the maximum density around the event horizon starts to catch up to the maximum density around the damage radius. At around 900 seconds, the maximum density starts to increase as the maximum density of the event horizon increases. The diminishing of high-frequency oscillations as time progresses is, again, due to the convergence of the disaster mobility model not found in Random Walk.

Figure 4 shows the variance of node density as the sim-




Figure 5: Clustering coefficient

ulation progresses in time. The variance of node density clearly increases as events are triggered. This is because many nodes have either a fairly small node density (if they are being partitioned or close to partition in the graph), or a high density (if they are clustered at the event horizon or damage radius). It is interesting to note that the high-frequency oscillations do not seem to diminish as time progresses. This is likely due to both the Random Walk civilians and the ambulances oscillating between hospitals and events.

Figure 5 shows the clustering coefficient in the network as the simulation progresses in time. The clustering coefficient of a node i is defined as:

$$C_i = \frac{2|\{e_{jk}\}|}{k_i(k_i - 1)} : v_j, v_k \in N_i, e_{jk} \in E,$$

where N_i are the neighbors of i , E is the set of edges in the graph, and k_i is the degree of node i [15] (this definition assumes an undirected graph). This gives a general indication of how well a node's neighbors know each other, which in turn gives insight into how clustered the network is. We define the clustering coefficient of a node with degree less than 2 to be 0. The average graph clustering coefficient increases sharply in response to the events. This is again due to clustering around the event horizon and damage radius for each event. As before, the diminishing of high-frequency oscillations is apparent.

At any given time, a graph is either partitioned or not. If it is, we say it has a "partition value" of 1. If not, it has a "partition value" of 0. For each data point, we have averaged the partition value of each of the 10 simulations. Figure 6 shows that the average partition value in the disaster mo-




Figure 6: Network partitioning

bility model increases as a result of the events. In fact, due to the high intensity of the events, after around 200 seconds the network is always partitioned in the disaster mobility model. This is because the nodes at the events are partitioned from the rest of the network, since the event horizons are generally out of their communications range. The disaster mobility model consistently has a partition value higher or equal to that of Random Walk, indicating a more fragile network.

These results show that our mobility model produces a topology much different than that of the popular Random Walk model. The vast difference between the topologies indicate that it is not sufficient to use Random Walk as a mobility model for disaster recovery networks.

7. CONCLUSIONS AND FUTURE DIRECTIONS

This paper presents a generic event- & role-based mobility paradigm used to characterize movement patterns of objects in response to environmental events. We specifically concentrate on applying this to disaster recovery scenarios, where current mobility models fail to realistically represent objects. To accurately characterize the movement of objects in response to one or more disaster events, we have developed a gravity-based model in which events emit forces that attract or repel objects depending on the object's role. Using simplified laws of physics, it is straightforward to calculate the velocity vector of an object, even in the presents of multiple events.

Our disaster mobility model has been fully implemented in simulation form and was used to generate ns2 mobility trace files. The resulting topology of our disaster mobility model had a higher average node density, maximum node density, variance of node density, and clustering coefficient. This is due to the grouping of nodes at or near the event horizons and near the damage radius of events. Furthermore, the partitioned value was consistently higher with our disaster mobility model, indicating the network was partitioned more often.

Our event- & role-based disaster mobility paradigm realistically captures objects' responses to disaster events. Furthermore, our simulation results show that the topological characteristics of the network drastically differ from that of Random Walk. As future work, we plan to perform studies on actual disaster scenarios to develop a rich set of role-based

rules and further refine our low-level gravitational model. Furthermore, we plan to use our disaster mobility model to understand the effects it has on routing protocols such as AODV [12] and DSR [8], as well as explore security concerns with a role-based system. This will most likely lead to the development of new DTN-style disaster routing protocols specifically tuned for disaster recovery networks.

8. REFERENCES

- [1] C. Bettstetter. Smooth is better than sharp: A random mobility model for simulation of wireless networks. In *ACM MSWiM*, 2001.
- [2] C. Bettstetter, H. Hartenstein, and X. Perez-Costa. Stochastic properties of the random waypoint mobility model. In *ACM MSWiM*, 2002.
- [3] C. Bettstetter and C. Wagner. The spatial node distribution of the random waypoint mobility model. In *German Workshop on Mobile Ad hoc Networks (WMAN)*, 2002.
- [4] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. In *Wireless Communications & Mobile Computing (WCMC)*, 2002.
- [5] V. Davies. Evaluating mobility models within an ad hoc network. Master's thesis, Colorado School of Mines, 2000.
- [6] R. Guerin. Channel occupancy time distribution in a cellular radio system. *IEEE Transactions on Vehicular Technology*, 1987.
- [7] A. Jardosh, E. Belding-Royer, K. Almeroth, and S. Suri. Towards realistic mobility models for mobile ad hoc networks. In *ACM MobiCom*, 2003.
- [8] D. B. Johnson and D. A. Maltz. *Mobile Computing*, chapter Dynamic source routing in ad hoc wireless networks, pages 153–181. Kluwer Academic Publishers, Feb. 1996.
- [9] B. Liang and Z. J. Haas. Predictive distance-based mobility management for PCS networks. In *INFOCOM*, 1999.
- [10] ns2 Network Simulator. <http://www.isi.edu/nsnam/ns/>.
- [11] M. R. Pearlman, Z. J. Haas, P. Sholander, and S. S. Tabrizi. On the impact of alternate path routing for load balancing in mobile ad hoc networks. In *ACM MobiHoc*, 2000.
- [12] C. E. Perkins and E. M. Royer. Ad-hoc on-demand distance vector routing. In *The Second IEEE Workshop on Mobile Computing Systems and Applications*, Feb. 1999.
- [13] G. Resta and P. Santi. An analysis of the node spatial distribution of the random waypoint mobility model for ad hoc networks. In *ACM Workshop on Principles of Mobile Computing (POMC)*, 2002.
- [14] M. Rossi, L. Badia, N. Bui, and M. Zorzi. On group mobility patterns and their exploitation to logically aggregate terminals in wireless networks. In *IEEE VTC*, 2005.
- [15] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, June 1998.