

The Application of Computational Models for the Simulation of Large-Scale Evacuations following Infrastructure Failures and Terrorist Incidents

Chris. W. JOHNSON¹,

*Department of Computing Science, University of Glasgow, Scotland.
<http://www.dcs.gla.ac.uk/~johnson>*

Abstract. Over the last five years, we have developed a range of stochastic techniques for simulating the evacuation of public buildings. These have included hospitals, entertainment complexes, a sports stadium, office buildings and concourse areas in an underground train system. This paper describes how simulations are developed from an initial risk assessment, through the integration of 3D architectural models and simulations of crowd behavior to the validation of results against 'live' evacuation drills and information about previous evacuations. The motivation behind this work is to support a form of 'mitigation engineering'. Evacuation simulations help us to reduce the consequences associated with a broad range of adverse events including fires, structural collapses and terrorist actions. This is especially important given the failure of risk assessment techniques to anticipate many different hazards that threaten critical infrastructures.

Keywords: Safety; Security; Evacuation; Terrorism; Infrastructure Failure.

Introduction

This paper provides an overview of recent research into the simulated evacuation of buildings that include hospitals, entertainment complexes, sports stadiums, office buildings and elements of the transportation infrastructures. An important motivation has been the idea that we cannot accurately predict all possible threats that face the general public. In consequence, we must devise emergency response strategies that mitigate adverse consequences without making undue assumptions about the nature of particular threats, which include but are not limited to natural disasters and terrorist attacks.

¹ Corresponding author: Prof. Chris. W. Johnson, Department of Computing Science, University of Glasgow, Glasgow, Scotland, UK, G12 8RZ. Email: Johnson@dcg.gla.ac.uk. The simulations described in this paper were developed as collaborative projects with a number of students over the last five years. Individual acknowledgements are grouped at the end.

1. Overview: The Move from Risk to Consequence Assessment

Risk assessment techniques focus finite development resources by considering the likelihood and consequences of particular hazards. We can formulate this in terms of a decision theoretic approach over vectors (*likelihood_0, utility_0; likelihood_1, utility_1; ...; likelihood_n, utility_n*). Each vector described the risks that are associated with the various scenarios that lead to particular hazards. For example, the risk associated with a fire could be calculated by summing the products of the likelihood and consequences (negative utility) of hazard scenarios including a match being dropped on the ground, electromagnetic discharge from a mobile phone igniting volatile gases and so on:

$$Risk_evaluation = \sum_0^n (likelihood_i \times utility_i)$$

Precise values for likelihood and consequence depend on the context that is being considered. For example, the risks of fires caused by electrostatic discharges from a mobile telephone hazard are very different in a hydrocarbon cracking unit from office environments. It is for this reason that techniques including fault trees, FMECA and cause consequence analysis have been developed to help analysts identify those hazard scenarios that should be considered within any particular instantiation of the risk formula given above.

This risk-based approach has been enshrined in national and international legislation. For example, most of the UK Health and Safety regulations include an expectation that employers will reduce risks until they are As Low As Reasonable Practicable (ALARP). Similarly, the Senate has urged the Secretary of Homeland Security to promote the National Fire Protection Association standard on Disaster/Emergency Management and Business Continuity. This requires that the owners and managers of public buildings conduct 'hazard identification and risk assessment' [17]. European Directives, 89/391/EEC and 89/654/EEC, describe minimum standards that should be enforced by legislation in each member state. The UK Fire Precautions (Workplace) Regulations were amended in 1999 to meet these directives by introducing a risk-based approach to fire regulations. Building owners and managers must take precautions that are appropriate to the likelihood and consequences of any hazard.

Risk-based techniques work very well for the analysis and design of hardware systems. The likelihood of particular hazard scenarios can be calculated in terms of failure probabilities using field data to calibrate bench tests. The likelihood of failure can be calculated in terms of the number of observed failures per operating hour. Structural decomposition can be used to assess the likelihood that a system will fail by considering the reliability of its various sub-components. Unfortunately, a number of problems limit the application of risk-based approaches. It can be difficult to quantify the likelihood of complex hazard scenarios. For example, to assess the likelihood of a fire being caused by a build-up of combustible materials in a railway station concourse, we must consider cleaning schedules for public areas, machine rooms and storage units. We must also consider the monitoring that is necessary to ensure that staff implement these schedules. In addition, we must also consider the likelihood of various ignition sources, including normal operating equipment, smoking materials etc. We may also have to consider contingent probabilities, where for instance the knowledge that a building company has broken regulations by leaving combustible materials in a station

area would justify an increase in the probability that workers will violate regulations involving tools, such as welding equipment, that provide potential ignition sources.

It is also difficult to assess the consequences of hazard scenarios. For example, the operators of the London Underground system had recognized the difficulty of entirely eliminating the possibility that a fire would occur on their premises. Instead, they trained their staff to detect fires and to extinguish them as soon as possible. For several years, this policy proved to be successful. There was a large number of relatively minor fires across the network. However, the fatalities from the Kings Cross tragedy illustrate the limitations of this approach. It is for this reason that many Health and Safety regulations embody the concept of 'worst plausible consequence'. Risk assessments cannot assume that a dropped match would result in a minor fire. However, it can be difficult to establish agreement about worst-case scenarios. These are often only revealed in the aftermath of major incidents. For example, the Kings Cross fire forced emergency personnel to look again at the possible consequences of minor fires causing flash-overs that lead to multiple fatalities.

Risk assessments must consider a broad range of different hazards, captured by different vectors in the previous formulation. Not only must we consider the likelihood and consequences of fires on the Underground, we must also consider the potential risks of terrorist attacks or a structural collapse [10]. For any particular hazard, how can we be sure that we have considered an adequate range of scenarios? In the previous example, we must consider many different combinations of fuels and ignition sources in order to account for the risk of fire. Information about previous fires in similar applications can inform this analysis, such as the Glasgow Underground, the Paris Metro or the New York system. However, it is often difficult to know whether information from other applications is relevant. For instance, the Glasgow Underground railway is far smaller than its London counterpart. There are fewer services, less passengers and the network topology is based on a simple loop. The Glasgow system does not deal with mixed over and under ground operations. Hence, information about potential causes of fire will be of limited value when applied to larger heterogeneous Metro systems.

The following paper accepts the difficulties that engineers and managers face in using conventional forms of risk assessment to support the design and operation of complex infrastructures. Rather than focus on techniques to produce particular likelihood or consequence metrics, the intention is to help engineers and managers develop emergency plans that reduce the adverse effects associated with a broader range of hazards. We call this *mitigation engineering*. It is impossible to guarantee that every possible risk has been considered. Systems should, therefore, be designed for resilience against many different hazards in an uncertain future.

This paper focuses on the evacuation of large, public buildings. This decision is justified by the impact of recent terrorist actions, such as the bombing of the London Underground, and the plethora of false alarms that follow such attacks. Fires, such as the destruction of the Station Night Club in Rhode Island, have also increased concern over emergency evacuations. Natural disasters, including Hurricane Katrina, and structural collapses, such as the damage to Terminal 2E at Paris' Charles de Gaulle airport, also guide the selection of case studies. We have chosen to focus on evacuation technologies because they provide an example of mitigation engineering. In other words, they provide means of reducing the consequences associated with a broad range of hazard scenarios. A further aim is to improve upon the subjective

walkthroughs that are relied upon for evacuation planning. These informal techniques have been widely criticized in the aftermath of major fires [9].

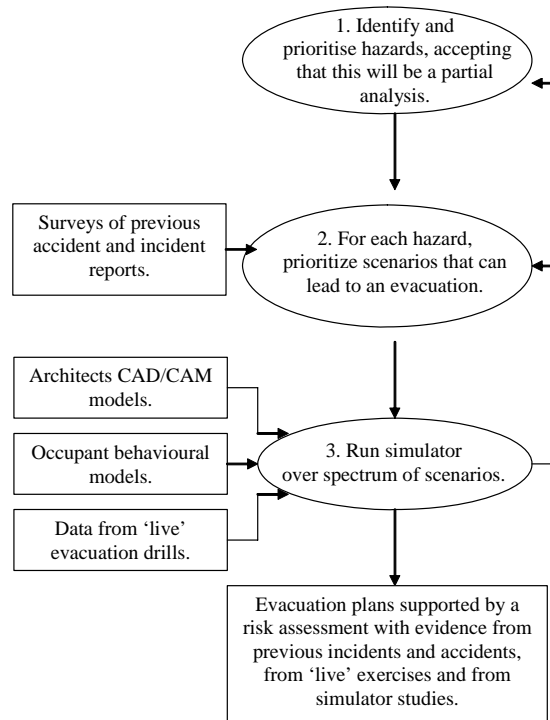


Figure 1: Overview of the Approach and Structure of the Paper

Figure 1 provides an overview of the proposed approach. An initial risk assessment can be used to identify the high-level hazards that affect large public buildings. These include fires, as well as crowd control problems, structural collapses and terrorist actions. The second stage helps to identify a range of scenarios that might lead to an evacuation following these hazards. It is important to stress that both of these initial stages provide indicative insights, they will not identify all possible scenarios for all possible risks. However, as mentioned above, by focusing on evacuation planning it may be possible to mitigate a broader range of potential risk scenarios than those that are explicitly considered during these initial analyses. The third stage develops and runs interactive evacuation simulations for the building and occupant population being considered. Subsequent sections of this paper will describe a suite of tools that automatically derive these simulations from the CAD/CAM files used by architects. This reduces the costs associated with simulation and also opens the potential to run evacuation simulations before a building is constructed [4]. This simulation stage also relies upon behavioral models for building occupants. For example, aggressive individuals often respond quite differently emergency evacuations when compared to large family groups [1].

The approach advocated in Figure 1 is iterative. Simulations can be shown to many different stakeholders, including building occupants and emergency personnel. These consultations often yield large numbers of additional hazards and evacuation

scenarios that must be integrated with existing risk assessments from the first and second stages. Similarly, annual or monthly evacuation exercises can yield further insights that must be incorporated into the evacuation planning process.

2. Concerns over the use of 'Live' Drills and Evacuation Exercises

The managers and owners of large public buildings currently rely on drills to help validate emergency evacuation plans. It can be difficult to ensure that these exercises recreate the range of conditions that hold during an actual emergency. There are obvious problems in reproducing the confusion, noise, smoke and structural degradation that characterize many incidents. Some organizations hold drills at the same time each week or month so that co-workers become habituated to these exercises, scheduling meetings to be out of the building. The costs and complexity associated with large-scale evacuation drills often dissuades organizations from considering a range of hazard scenarios. Exercises become piecemeal efforts that seldom involve the police or fire services.

There are also dangers in drills for individuals with pre-existing cardiovascular conditions. The risks of injury will often dissuade rail operating companies and airlines from hosting evacuation exercises with members of the public. In such circumstances, other members of staff 'act' as customers. However, this often results in drills that do not accurately recreate the diverse groups who must be evacuated during a real emergency. By recruiting groups of co-workers, drills often fail to recreate the behaviors of family groups, which tend to evacuate or remain together during an emergency. The Federal Emergency Management Agency has argued that the stronger the bond between group members, the more likely it is that one member will put their own life at risk to protect another group member. There are particular problems in organizing drills that include the very young and the very old.

Further problems limit evacuation exercises in public buildings. Few assumptions can be made about the prior knowledge of many building occupants. Customers in cinemas, theatres and clubs change from hour to hour and day to day. Hence they may not be aware of emergency exits or of evacuation procedures. Drills, therefore, tend to focus on the key role played by a small number of staff. The owners and managers of public buildings must also consider crowd control issues, especially in entertainment venues. For example, it can take more than an hour to partially evacuate stadiums that accommodate more than 60,000 seated spectators [8]. The costs of recruiting participants in such drills are prohibitive.

Particular problems restrict the use of evacuation drills in healthcare settings. Several US states have prohibited or restricted the involvement of 'real' patients. Even exercises that involve actors rather than patients can disrupt 24/7 healthcare provision. Exercises often have a limited validity in clinical environments. For example, the evacuation of an operating theatre depends on a close cooperation between the nursing, surgical and anesthesiology teams that must cooperate to stabilize the patient and move them from the area at risk. However, these teams often work in rotations. There are few guarantees that different teams will achieve the same levels of performance in successive drills. Changes in team members, the complexity of the surgery and technologies all affect evacuation performance.

Safety concerns are compounded by security issues in some environments. For instance, banks and other financial institutions face external threats that might

exploit vulnerabilities created during evacuation drills. Security concerns can be strongly linked to the evacuation of key personnel during drills.

3. Overview of Interactive Evacuation Software

Simulation software can overcome many of the limitations mentioned above. Building managers can use them to consider a broad range of potential evacuation scenarios. For example, evacuation tools can be configured to calculate probable evacuation times for different occupancy levels within a public building. It is important to stress that the insights obtained from these applications must be validated and calibrated against the results obtained from 'live' drills. Simulation tools can also be used to explore the possible consequences of changes in the layout or function of a structure before they are implemented.

There are numerous disadvantages with evacuation simulators. They can be expensive to build, especially for legacy buildings where there may not be any computer-generated plans to assist in software development. It can also be very difficult to validate predictions. If the models are too pessimistic then building managers and operators will invest in safety improvements that cannot otherwise be justified. If they fail to anticipate potential evacuation problems then they can instill overconfidence that will be undermined by any emergency. It is for this reason that software tools can only be seen as extensions to drills. Their use must also be informed by reference to reports on previous evacuations from other, similar structures.

The success of an evacuation simulator is determined by its model of human behavior. There are many different aspects to be considered. For example, building occupants often try to establish the credibility of an alarm before starting to move towards a place of safety [2]. Simulations mimic these findings by introducing a fixed delay before occupants move in each run. More elaborate systems may also model the perceived threat posed by the alarm, the occupant's preoccupation with the task to hand, familiarity with evacuation procedures from previous drills etc. Simulators must also consider social and team factors. For instance, Tong and Carter [16] describe "flocking", which attracts more people into areas that are already crowded.

Evacuation simulators must also account for different physiologies. This determines the different speeds at which, for example young and old, will travel through the building during an evacuation. It is also important to stress that the physiology of building occupants need not be normally distributed. Hospitals provide the most obvious examples where some patients are confined to beds or wheel chairs while others can walk. The proportions in each category vary from ward to ward. Similarly, the average walking speed and cardio-vascular capacity of building occupants can be skewed in sports centers, prisons and nightclubs. It is for this reason that physiological profiles are often constructed to help tailor simulations to particular environments. For example, timings can be taken during evacuation drills to help anticipate mean walking speeds between different areas of a building. These measures must account for the ways in which velocity varies with crowd density. Individuals slow their pace well before reach a crowded area. It is also important to gather empirical data for the time taken to negotiate stairs under a variety of emergency conditions.

It is relatively easy to compile physiological information about occupants and then use that data in an evacuation simulation. For instance, if we observe that 60% of a

building's population are capable of moving at 1.4 ms^{-1} over flat ground then we can develop code to ensure that this proportion of the simulated occupants will move at that speed every time it is run [15]. Similar techniques can also be used to calibrate the cognitive profiles that are used by evacuation simulators. For example, psychometric questionnaires may reveal that 30% of the people in an area of a building exhibit a low level of aggression, 40% reveal medium levels and 30% exhibit high levels of aggression. Precise distributions vary from population to population. For example, the levels of aggression observed in a hospital Accident and Emergency departments can be very different from those in other in-patient departments. The psychological and physiological profile of building occupants also varies over time with changes in the occupancy of a public building. For instance, there are considerable additional difficulties in evacuating an Accident and Emergency department on Friday and Saturday evenings; the patient profile is very different from that of a weekday morning. These differences can be captured by studying the occupants of the building, as mentioned above, and then altering the distributions. For example, during a weekday morning the proportion of highly aggressive patients (and staff) may fall from 30% down to 10%. This reflects the lower likelihood of having to deal with aggressive patients outside the weekend.

Many simulators use Monte Carlo techniques to model the small differences in cognitive and physiological profiles that can be observed in the occupants of many buildings. Random numbers are compared with probability distributions to determine the population profile for any particular run of a simulator. In the previous example, if we were simulating the patients in an Accident and Emergency department during the weekend where 30% of the population was found to exhibit high levels of aggression and the random number fell in the range from 1 to 30 then the simulated occupant would exhibit behaviors associated with high levels of aggression. Correspondingly, if the number fell above this range then they would be classified as exhibiting behaviors associated with moderate or low levels of aggression depending on the value of the random number. These techniques provide a link between the psychological profile of the occupants in a particular building and human factors studies of evacuation behavior where, for example, levels of aggression have been shown to provide good indicators of mean evacuation time [12]. The use of Monte Carlo techniques helps to ensure that the precise composition of a population will change between different runs of the simulation because the random numbers will also change. However, the use of probability distributions ensures that over repeated runs, the distribution of the population between each category will converge on the proportions identified by the psychometric questionnaires and observational studies. If we were to increase the proportion of aggressive individuals in our simulated population then this would increase the likelihood that any individual random number would fall within the associated range. Hence it would be more likely that an individual in the model would exhibit the associated behaviors. The use of non-deterministic techniques helps to ensure that we do not always see high levels of aggressive behavior in every simulation of an Accident and Emergency Department. However, the use of probability distributions ensures that over repeated runs, the simulations reflect behaviors that are associated with the psychological and physiological profiles of individual occupants.

The same Monte Carlo techniques can also be used to ensure that individuals in a model do not always behave in the same way during successive runs. In this case, probabilities are associated with different courses of behavior. Less aggressive individuals are more likely to wait for other people to clear away before they move

towards an exit. More aggressive individuals are more likely to move forward before the exit is cleared [12]. However, these behaviors are not deterministic. In other words, aggressive individuals do not always push ahead. Less aggressive individuals do not always wait. Random numbers are, therefore, generated and compared against probability distributions describing the likely behaviors for individuals in each of the population groups. This ensures that building occupants do not always follow the same course of action during each run of the simulation. They are, however, more likely to perform those actions that are considered to be most probable for their group in a particular scenario during an evacuation. The probability of particular behaviors can be directly informed by previous incident reports and by the observations derived from evacuation exercises. This further strengthens the links between software simulation, previous incidents and evacuation drills illustrated in Figure 1.

It can be difficult to identify the associations between particular evacuation behaviors and psychological or physiological attributes. Studies such as those by Latman [12] linking aggression to mean evacuation times are relatively rare. Further problems stem from the way in which Monte Carlo techniques mimic the observable behavior of building occupants. There is no pretence that they capture the higher-level problem solving techniques that individuals employ during real evacuations. Further limitations can be addressed by exploiting more complex mathematical techniques, such as Markov chains or Bayesian networks, to capture changes in occupant behaviors over time. However, it is important to recognize that *simulations will only ever provide approximations for live exercises and these only provide approximations of the behaviors observed under emergency conditions*. Exact estimates of individual evacuation times are of less importance than order of magnitude figures which must then be validated by drills.

4. Simulating the Evacuation of Large Public Buildings

The Glasgow Evacuation Simulator (GES) uses Monte Carlo techniques to support the iterative approach illustrated in Figure 1. Figure 2 shows the GES being used to model evacuations from a public, auditorium complex in Glasgow, Scotland. Users can vary the occupancy levels in the building. They can also interactively open and close emergency exits as a simulation progresses to model the effects of damage to the building or intervention from the emergency services. It is also possible to specify that occupants will exhibit 'model behavior' in which they are likely to use the nearest available emergency exit. Alternatively, the software can be used to simulate scenarios in which most occupants retrace their steps back towards the main entrance for the building. This is important because fire safety regulations assume that drills help occupants learn the location of emergency exits. The requirement to conduct regular exercises is based partly on the assumption that occupants will practice using these escape routes. However, many people ignore emergency exits during emergency evacuations. This was a feature of the Lowenbraukeller and Station Night Club fires [9]. We conducted a series of surveys and interviews in an attempt to persuade building occupants that they should use the emergency exits illustrated in Figure 2. However, we were unable to convince 25% of the occupants. They cited a wide variety of reasons for preferring to retrace their steps back along the route that they had used to enter the building. These included an unfounded concern that the building managers

had barred fire exits in order to secure the building. Others argued that they had never used the emergency exists and did not know where they would lead.




Figure 2: User Interface to the Glasgow Evacuation Simulator (GES)

Figure 3 illustrates some of the results that were obtained from the GES tool. The graphs show the time taken to clear the building when an emergency stairwell is blocked. The top line shows mean evacuation times under different occupancy levels when individuals are likely to retrace their route into the building. The lower line provides the same information for 'model' evacuations in which each occupant attempts to exit by the nearest available route. The difference between the 'model' and 'normal' mean evacuation times is much greater than when any of the other emergency stairwells are blocked. Hence, considerable efforts should be made to ensure that building occupants use these routes rather than retracing their steps if they are to benefit from the time savings indicated in Figure 3.




Figure 3: Graphing Mean Evacuation Times when an Exit is blocked

The insights provided by graphs such as those shown in Figure 3 should not be underestimated. Both the graphs and the three dimensional simulations can inform the emergency procedures developed by building managers. In particular, by illustrating the increasing delays that are created when this particular stairwell is blocked. The graph reiterates the importance of marshals and fire safety officers guiding occupants to the remaining stairwells. These results can also direct intervention by the emergency services, for example, by illustrating the relative importance of this exit route. The interactive nature of the GES tool also enables its use during exercises, fire service personnel can simulate what might happen if other emergency exits are progressively lost during a structural collapse or fire. As we have seen, there are significant financial, organizational and ethical barriers to the evaluation of these different scenarios using live drills.

5. Simulation in Planning and Design

Previous simulators have forced managers and owners to develop bespoke models of their buildings. This often involves significant costs as developers translate paper plans into software simulations. These models are complex because they must capture the layout of the building and also encode details about its function and operation. This is important if software engineers are to ensure that the simulated occupants of a building do not escape by walking through the walls of a model. Similarly, programmers must encode the dimensions of doors, corridors and stairwells in order to ensure that any simulation accurately captures the capacity of these different areas. Models may also distinguish between different construction materials, including fire and smoke resistant doors. The costs of developing the graphical representations and

structural models have acted as an important barrier to the application of computer-based simulation. Compromises have also been made to lower the associated overheads. For example, some tools only support two dimensional models. Others simplify the movement between floors by ignoring stairwells. They focus in on relatively small subsections of larger buildings and hence have only a limited application.

Variants of the GES tool have addressed these limitations by reusing existing 3D models from architects' design tools. Unlike many other simulators, there is no need to build specialized visualizations. This reduces costs and allows a tight integration between the simulator and the design of new structures. The ability to derive simulations from the files of tools such as AutoCAD enables us to simulate buildings that have yet to be constructed. We can model evacuations to assess the likely impact of small changes in the layout of a proposed building. There are, however, a number of outstanding problems. In particular, most architectural design tools use proprietary file formats to link 3D projections with semantic information. These formats show that individual lines can be combined to form more complex objects, such as doors, corridors or stairwells. We cannot access these detailed formats without greater support from the developers of the CAD/CAM tools. However, we can incorporate the 3D drawings. Users must, therefore, still develop data structures to represent these key areas of a building, even though we can directly re-use the 3D visualizations.




Figure 4: State of Building during Initial Evacuation Simulation

Figure 4 provides an example of simulations being used before construction is completed. The building is relatively complex because it will extend and be connected to the row on the left of the image. It will be possible for occupants to exit through the existing building or by using stairwells within the new structure. The photograph illustrates the state of the construction at the point when the simulator was used to make initial predictions about potential evacuation times. Figure 5 shows the evacuation simulator that was developed. As before, it is possible to block and

unblock exit routes during the simulation. The four panels at the bottom right of Figure 5 provide an overhead view of the stairwells. The top right panel provides an alternate projection between floors. These areas form a particular focus for the evacuation model because they are potential bottlenecks.




Figure 5: Simulation of Occupant Egress and Fire Crew Ingress before Completion of Building

Many previous simulators have focused exclusively on the egress of building occupants. The attacks on the World Trade Center illustrated the limitations with this approach; by focusing on egress we may ignore the potential risks to emergency service personnel who must secure the occupants of damaged buildings [9]. This is one of the reasons why the previous model was extended to model the stairwells in considerable detail. Emergency service personnel must often battle against the flow of building occupants in order, for example, to locate a fire. The development of the simulator illustrated in Figures 4 and 5 helped to reveal considerable variations on the speed of ingress of emergency personnel depending on the time taken to initiate the evacuation. If occupants began to leave the building within a few second of the alarm being issued then the stairwells would be clear for the arrival of the emergency personnel; most occupants would have left by the time that they arrived. However, if the initial delay after an alarm stretched to anything more than one minute then emergency personnel were significantly delayed in reaching the fire as people continued to stream out of the simulated stairwells.

The system illustrated in the previous diagrams provided further insights into the close relationship between the evacuation of the proposed extension and the existing building. Occupants could choose to exit horizontally from the proposed structure onto the corresponding floor of the existing building before using one of its stairwells. This raised a number of issues. If occupants were encouraged to use this route then it cleared the stairs in the new building for the arrival of the emergency services. However, it could also create further bottlenecks on the landings in the

existing building. The present occupants of the existing building were joined by a number of additional people from the proposed building. Such observations show that simulators raise as many questions as they answer. It was decided to address these potential issues at a later stage in construction when it would be possible to gather empirical observations of occupant behavior to back-up the insights from the software models.

The simulation of the proposed building illustrates further technical innovations. In the previous model of the auditorium complex, individual occupants are controlled by an executive function that periodically cycles through internal data structures to update the position and planned path of each simulated person. In contrast, each occupant in the simulation from Figure 5 is represented as a separate thread or process. They continually and asynchronously update their position. Other simulated occupants will react to the changes in the path of their co-workers through various forms of message passing. This approach offers significant benefits in terms of the software engineering; it is possible to create more complex cognitive and physiological models without causing widespread changes to the rest of the code. However, there are computational overheads that make it difficult to sustain this approach for simulations involving tens of thousands of occupants.

6. Sporting Venues

Recent years have seen a number of adverse events involving sporting venues. These include the collapse of barriers around stairwells at Ibrox Stadium in Glasgow (1971), the Bradford Stadium fire (1985), crowd disturbances compounded by decaying infrastructure at the Heysel stadium, Belgium (1985), crush injuries against steel fencing at Hillsborough in Sheffield (1989), a stand collapse and crush at the Armand Césari Stadium on Corsica (1992), panic following police use of tear gas at the Accra Sports Stadium (2001), overcrowding and a subsequent stampede at Ellis Park in Johannesburg (2001). These incidents had very different causes. It is, therefore, appropriate to apply 'mitigation engineering' techniques to assist evacuation following many different adverse events with stadium buildings. It does not automatically follow that computer-generated models are an appropriate means for conducting such an analysis. Still has argued "...tests can only be carried out by using real spectators. Bank computers and other computer technologies can be tested by putting through figures, data and printout. But the system can only be tested by putting through people. So in the very testing of whether the system may cause danger, danger may be caused...The laws of crowd dynamics have to include the fact that people do not follow the laws of physics, they have a choice in their direction, have no conservation of momentum and can stop and start at will. They cannot be reduced to the equations that are appropriate for the movement of ball bearings through viscous fluids" [14].

These are valid criticisms. However, they miss the point of the integrated approach illustrated in Figure 1. The intention behind our use of computer simulations is to maximize the insights that can be obtained from 'live' studies [6]. Although we cannot eliminate the risks involved in these large-scale live drills, it is important to minimize any potential hazards. Simulations can be used to rehearse evacuation drills. This is essential for stadium exercises where several thousand people are recruited to test emergency procedures. The results from exercises can then validate initial predictions from software tools. There are some scenarios that cannot easily or

accurately be recreated using live drills. It is equally difficult to validate actors' responses to simulated baton charges or structural collapses as it is to validate human behavioral models in software simulations.



Figure 6: Photograph and Simulated Model of Sports Stadium

Figure 6 illustrates early stages in the development of an evacuation simulation for a large sports stadium. The software is intended to exploit detailed observations from previous incidents, as well as insights from live drills and the expert advice from police safety officers. A particular issue here is that it is impossible to conduct full scale exercises in which all stands are filled and then evacuated under a number of different conditions. In addition to the usual ethical and logistic problems with drills, there are significant concerns for the local population and transport infrastructure. In consequence, most stadium managers conduct a limited number of partial exercises involving restricted areas of a venue, typically focusing on a single stand. This makes it difficult to simulate the redeployment of emergency personnel from one area of a stadium to address an incident in another stand. Single-stand drills cannot easily be used to predict the additional overheads associated with multiple problems occurring simultaneously in different areas of a stadium. This is a common feature of several scenarios for terrorist action [11]. By confining live exercises to isolated areas in an arena, it can be difficult to reproduce the problems that arise when different groups compete for finite amounts of space to set up triage areas, communications posts, marshalling stations and so on.

In contrast, simulation software can be extended from the main auditorium to stairwells, passageways, concourses and car parks surrounding many venues. The insights from limited drills can be 'scaled up' to illustrate the combinatorial logistics that characterizes multiple incidents. It is relatively easy to establish game-playing scenarios for what might happen when communications infrastructures break down or when key members of staff become side-tracked by detailed concerns. These have all been features of previous disasters involved mass crowds at sports stadiums.

There are of course a number of problems that complicate the application of these simulators. Previous sections have described the use of Monte Carlo techniques to introduce non-determinism into individual behavior. These approaches can also simulate group responses to emergency situations. For example, probability distributions and random numbers can be used to identify specific objectives for collections of spectators within particular areas of the stadium. A second, lower-level iteration can then be used to apply Monte Carlo techniques as a means of simulating individual variations within the group. However, such approaches place heavy demands on the computational resources of most personal computers when applied to crowds of over 50,000. This is an important consideration given that we would like the tools to be used by the many different groups involved in emergency evacuations. Initial prototypes are, therefore, designed around a simplified form of the central executive used in the public auditorium system, introduced at the start of this paper. The intention is not to provide an 'accurate' model of group behavior but to provide an environment that provokes emergency personnel to consider a wider range of hazard scenarios and thereby maximize the benefits to be obtained from a limited number of live drills.

7. Hospital Evacuations

The difficulty of predicting the many different hazards that might trigger an evacuation together with the potential adverse consequences of those hazards makes it critical that hospital staff prepare detailed emergency procedures. As with sporting venues, this argument is reinforced by a number of previous incidents. For instance, there are approximately 2,500 major fires in Scots hospitals alone. In the United States, there are 3,500-4,000 fires involving multiple fatalities in nursing and assisted living homes per annum. No accurate records are kept for the number of incidents that lead to the deaths of single individuals. The focus on hospital evacuations is also justified by public concern following particular incidents [3]. In 2003, 30 patients died in a hospital fire in Belarus while another 10 died in a fire at the Greenwood Health Care Center in Connecticut, USA. In January 2004, the Rosepark Care Home fire in Uddingston killed ten patients and sparked a national debate on the safety of healthcare institutions in the UK.

The ethical and logistic problems that arise during mass evacuation drills from sports venues are more complex when considering emergency planning for healthcare organizations. It is also important to illustrate the scale and complexity of evacuation exercises in these environments. For example, a US hospital recently conducted a number of fire drills over a 6-week period. One scenario started when the tip of an electrosurgical pencil ignited a drape or cover [13]. Staff members rapidly removed the cover from the patient by throwing it on the floor. The 'simulated' fire was extinguished. At this point, however, the staff running the simulation intervened to inform them that the fire had spread. There was initial confusion; several adjacent rooms were evacuated at the same time causing temporary gridlock in the corridors. The drill simulated the movement of intubated patients using the operating room bed with a bag-valve mask. The exercise also required staff to move individuals with open incisions. Wounds were packed with sterile, saline-soaked laparotomy sponges and then covered with sterile drapes. The evacuation scenarios were also scripted to determine whether staff knew which items of equipment needed to be evacuated with

their patients. They had to collect enough instruments to close the incision even though the evacuation plans provided for sterile equipment to be available in the triage area.

These drills are necessary because staff must rehearse complex procedures including the ‘horizontal evacuation techniques’ advocated in UK hospitals. Patients are moved from a hazardous area to a place of safety on the same floor, for instance behind fire resistant doors and walls. Patients are only moved to other floors or out of a building as a ‘last resort’. The evacuation follows a predetermined plan in which staff must first locate the source of any hazard and then ensure that the proposed destination will keep them free from danger until the emergency services arrive. Staff must continue to ensure that there is a protected route from the place of safety to an exit from the building. Different classes of occupant raise different concerns during an evacuation. Some assessment may have to be made about whether the risk of moving the patient is greater than the risk posed by the fire or other hazard. Patients must be taken to a place of safety that does not impede the ingress of emergency personnel. This is important because there is a danger of injury as equipment and people move in to tackle a fire or similar hazard.

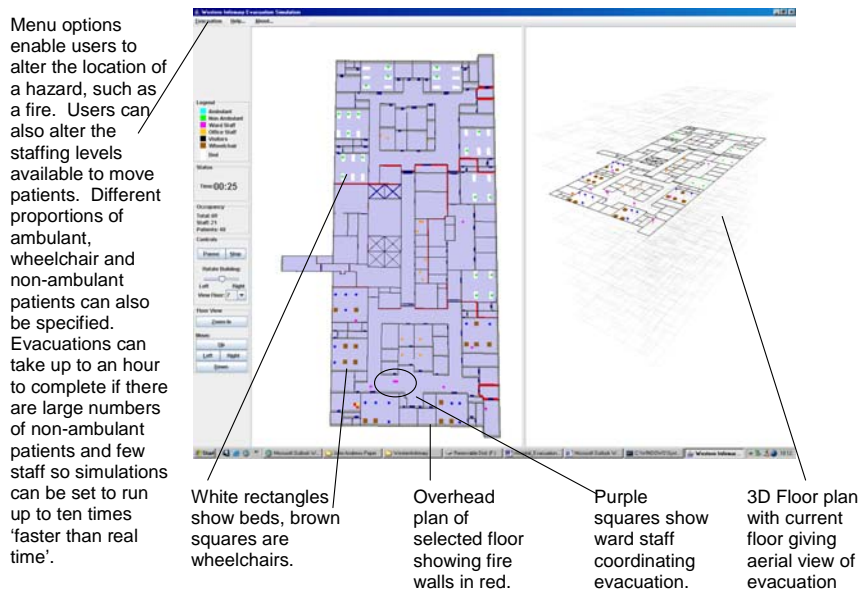


Figure 7: Glasgow-Hospital Evacuation Simulation Showing Staff and Patients

Figure 7 presents the interface to the Glasgow Hospital Evacuation Simulator (G-HES). As before, CAD/CAM projections can be incorporated into the tool. However, the use of Monte Carlo techniques must be balanced against the more prescribed behaviors associated with horizontal evacuation. In other words, cognitive differences between nurses are assumed to have less influence over their actions as they follow the various steps in the hospital evacuation plan. All patients in immediate danger are moved first. Next ambulatory patients and visitors are moved. Wheelchair patients are grouped together and then moved gradually to a place of safety. Finally, non-ambulatory patients will be moved. The implicit objective at each stage is to maximize the

number of people who can be evacuated in the shortest period of time. In addition to modeling these task priorities, it is important for the simulation to model a number of delays. For example, we conducted a number of empirical studies to assess the time that is necessary before a patient can be moved. This may involve transferring patients from fixed monitoring equipment to portable counterparts. It can also involve changes in the wiring or arrangement of medical devices including infusion pumps and breathing equipment. Further studies were conducted to determine the speed with which nurses can move patients who are confined to their beds or to wheelchairs. These studies were critical because they revealed considerable differences between different models of bed and chair within the same wards.

Table 1: Evacuation Times for Day Staff of 6 Nurses with 10 Runs for Each Patient Distribution

Number of Non-Ambulant Patients	Number of Ambulant Patients	Mean Evacuation time in seconds (Min:Sec)	Standard Deviation in seconds (Min:Sec)
30	0	2643 (44:03)	257 (4:17)
25	5	1749 (29:09)	205 (3:25)
20	10	1439 (23:59)	189 (3:09)
15	15	1105 (18:25)	86 (1:26)
10	20	801 (13:21)	75 (1:15)
5	25	707 (11:47)	64 (1:04)
0	30	470 (7:50)	54 (0:54)

Table 2: Evacuation Times for Night Staff of 3 Nurses with 10 Runs for Each Patient Distribution

Number of Non-Ambulant Patients	Number of Ambulant Patients	Mean Evacuation Time in Seconds (Min:Sec)	Standard Deviation in seconds (Min:Sec)
30	0	3445 (57:25)	363 (6:03)
25	5	2976 (49:36)	279 (4:39)
20	10	2703 (45:03)	253 (4:13)
15	15	2357 (39:17)	234 (3:54)
10	20	1991 (33:11)	226 (3:46)
5	25	1723 (28:43)	244 (4:04)
0	30	1343 (22:23)	227 (3:47)

Tables 1 and 2 summarize initial results that were obtained using the simulation system illustrated in Figure 7. The intention was to explore what might happen to evacuation times with different profiles of ambulant and non-ambulant patients under the given staffing regime within particular areas of the hospital. We, therefore, began to apply

the tool by examining ten separate runs for the current staffing level of six nurses faced with different proportions of ambulant and non-ambulant patients. The results are shown in Table 1. Table 2 illustrates the increased evacuation times associated with the reduced staffing levels during night shifts. These figures are illustrative. The simulations do not consider the additional complexity of rousing patients from sleep when they may be under additional sedation. Neither do they account for the additional fatigue that may be expected if a small number of staff work on evacuating patients for long periods of time. These factors could be modeled in the software. However, it is far harder to conduct nighttime validation exercises through live drills. It is unclear whether it would ever be possible or ethical to obtain staff participation to assess fatigue in an exercise involving non-ambulant patients where simulation results indicate it might take an hour or more.

8. The Evacuation of Transportation Infrastructure

Figure 8 presents the user interface to an evacuation simulator for a station on the Glasgow Underground. In contrast to other simulations where occupants must descend to escape from a building, passengers must ascend from the platform to reach the surface. The behavioral models must capture the confusion and disorientation that characterizes the evacuation of underground rail systems [7]. Simulators must also account for the physiological differences between potential passengers. They must model family groups and individuals, they must also account for the very young and the elderly. They must capture different physiologies; many passengers require staff assistance. It is also important to consider passengers who cannot be dissuaded from bringing personal affects and luggage over flights of stairs, escalators and turnstiles, as they evacuate the station illustrated in Figure 8.

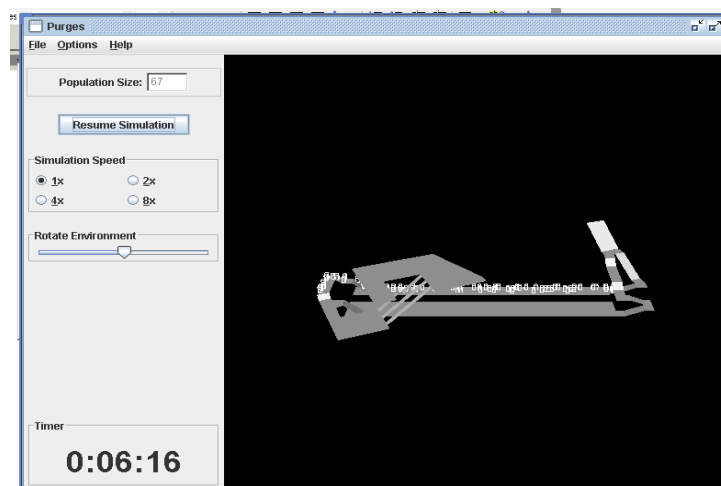


Figure 8: Evacuation of a Single Platform in an Underground Rail Transit System

It is particularly difficult to conduct live exercises for the evacuation of mass transit systems. Commercial and ethical factors prevent the involvement of paying

passengers during normal operating hours. Drills must usually be run at night when passenger services are suspended. This makes it difficult to reflect the broad population of travelers, especially young families.

Just as the previous section raised questions about the relationship between the planned extension and the existing building, the model of underground evacuation also raises important questions about the scope of the simulation. In this case, developers had to represent the flow of passengers into the station as well as their egress from the platform. In the other simulations, the occupancy of a building will not rise during an evacuation once we have considered the entry of emergency personnel. In the Underground, new passengers are brought into stations with the arrival of subsequent trains. This forces managers to consider the length of time to hold a train outside a station in order to allow the previous group of passengers to clear the main exit routes. It is both difficult and dangerous to run several different scenarios during live drills. However, the simulation software can be extended to model a range of different evacuation procedures so that subsequent trains are not delayed unnecessarily and the overall success of any emergency response is not jeopardized by overcrowding if too many passengers are allowed into a station.

The simulation illustrated in Figure 8 also embodies many of the aims of mitigation engineering, described in the opening sections of the paper. The evacuation models can consider a range of potential threats. These include fire and structural collapse. However, we are also considering the use of the system to plan the ingress of emergency personnel following a bioterrorism incident. The initial development of this scenario is based on the sarin gas attack by members of the Aum Shinrikyo religious group on 20th March 1995, against subway trains passing through Kasumigaseki and Nagatcho, Japan. In such circumstances, the need to provide emergency assistance and urgent medical evacuation must be balanced against the potential risks to fire and police crews as they enter the subway system. At a more immediate level, managers must make critical decisions about the containment of passengers who may not yet be contaminated as they approach stations where there is evidence of a previous attack.

9. Disease Control and Large Scale ‘Self-Evacuations’

Many simulators focus on individual buildings. For example, the previous section described a relatively contained bioterrorist attack on one section of a transportation system. However, the World Trade Center attacks of 2001 revealed the dangers of such a limited scope [9]. Public attention focused on the Twin Towers. Fatalities, injuries, building collapses and structural damage extended many blocks from the points of impact. Similarly, the events of the al-Aqsa Intifada have shown that determined groups can launch coordinated attacks throughout entire districts of major cities. The myriad of problems that complicated FEMA’s response to Hurricane Katrina also illustrate the importance of expanding the horizons for emergency planning.

We have, therefore, begun to consider the impact of mass evacuations. In the past, it would have been inconceivable to develop photo-realistic simulations of major population centers. However, a number of commercial and public research projects have developed three dimensional city models [5]. The focus has been on developing backgrounds for computer games or on providing tourists with information about

different attractions. These digital resources can be re-used in the same way that we have exploited the existing CAD/CAM models of individual buildings. Digital city models already provide semantic information about the function and population of various buildings using publicly accessible formats. This forms a strong contrast with the proprietary data structures that are embedded in many architects' tools.

This work is in its early stages. Particular problems relate to the lack of data on real evacuations or mass drills involving major conurbations. Information is available, for example about areas that are frequently evacuated in anticipation of major storm and hurricane damage. However, these areas often have established evacuation plans, dedicated lanes on interstates etc. It is unclear whether insights from these regions can be applied to other areas with less experience of mass evacuations. It can also be difficult to determine whether existing preparations would hold up under the new hazard scenarios that are being considered in the aftermath of Katrina.

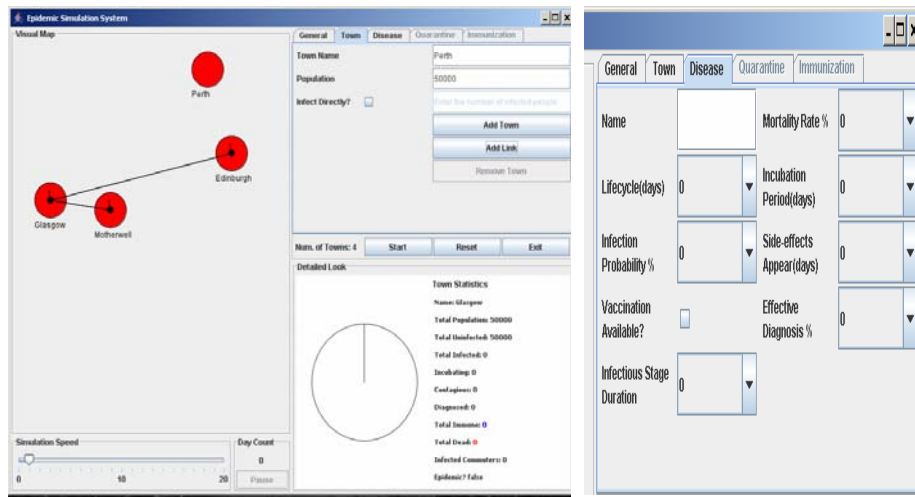


Figure 9: Considering the Interaction between Infection and the 'Self-Evacuation' of Conurbations

Figure 9 presents an alternative approach to the simulation of mass evacuations. This does not rely upon three dimensional models of urban districts. Instead, it provides a more abstract representation of the epidemiological implications of 'natural' and bioterrorist hazards. The screen on the left provides a high-level map of neighbouring towns. The red circles are pie charts that indicate the proportion of the total population affected by a disease. The links between the circles indicate population flows between these towns and cities. The image on the right presents the panel that is used to tailor the simulation to reflect the lifecycle and morbidity of particular diseases. It can also be used to assess the impact of different control measures including quarantine and immunization. These parameters can be tailored to relatively well known profiles, such as smallpox, or they can be used to analyze different scenarios that might emerge, for instance for different forms of avian influenza.

There are strong links between this simulation and the previous evacuation models. They all rely on a mixture of empirical information, for example physiological observations or disease characteristics. They also all exploit location information, ranging from building plans to data on the transport links between cities. They must

also capture elements of human behavior. This is particularly important for epidemiological simulators. We can specify the probability of transmission for particular diseases but this can only be used to make predictions if we know the number of contacts that an infected person has in a given period. Similarly, the distribution and spread of any disease is determined by the effectiveness of quarantine measures in limiting the understandable desire for individuals to 'self evacuate' areas involved in a bioterrorist incident. If an immunization programme is used in conjunction with quarantine measures, then it is important to assess the level of participation within a population.

The flexibility of the interface shown in Figure 9 partly reflects our emphasis on 'mitigation engineering'. The intention is to avoid undue constraints on the range of hazard scenarios that might be considered during emergency planning. The aim is also to develop robust strategies that might prove effective against many different potential threats. We do not know the likely morbidity of future epidemics and hence the planning must consider different scenarios. However, the generic and unconstrained nature of the interface in Figure 9 also reflects our wider ignorance about the public response to such events. The difficulty of organizing evacuation drills for individual buildings seems trivial compared to the problems of anticipating the public response to mass quarantine and immunization programmes.

10. Conclusions and Further Work

This paper has provided a high-level overview of several different evacuation simulators including hospitals, a sports stadium, an auditorium complex, a proposed office building and elements of the transportation infrastructure. We have described how these simulations were developed from an initial risk assessment, through the integration of 3D architectural models and simulations of human behavior through to the validation of our results against 'live' evacuation drills. The intention has been to support 'mitigation engineering'. Given that we cannot anticipate the many different hazards that threaten critical infrastructure, simulations help us to prepare emergency procedures to reduce the consequences associated with a broad range of adverse events including fires, structural collapses and terrorist actions.

Given the broad scope of this paper, we have only been able to introduce many issues. For example, we have not described the detailed communications mechanisms that teams of co-workers use to coordinate their response during an evacuation. Similarly, we have not considered the techniques that are used to integrate the simulation software into wider forms of training where emergency personnel engage in game playing. Citations have, however, been provided for further reference. Work is continuing to ensure that our computational techniques provide valid approximations for the behavior observed in live exercises and previous incidents in mass crowds in stadiums. Similarly, we are working on a range of application domains that pose particular challenges for the techniques that are described in this paper. In particular, we are using studies from a number of previous nightclub fires to develop simulations for a range of venues that are concerned about the particular problems of evacuating large numbers of young people. We are also working on simulators that model the evacuation from cruise ships, in particular, to address the problems that might arise in the immediate response to a terrorist attack in port or at sea. Although these applications are associated with relatively detailed hazard scenarios, our work

continues to focus on a resilient approach. Evacuation procedures must provide a flexible response to the wide range of potential threats that cannot easily be predicted in an uncertain world.

11. ACKNOWLEDGMENTS

Thanks are due to the teams of students who have contributed to the development of the simulators over several years. The initial evacuation of the auditorium complex was driven by J. Appleby, P. Cooper, A. Foss, S. Hailey and B. Jenks. J. Johnston and G. Mckinlay were responsible for the development of the 'new building' simulator. J. Ferguson simulated evacuation from the Underground station. P. LaValle, M. Trybilo, W. Munnoch, and S. Banks developed the evacuation of the sports stadium. F. Ashraf, J. Johnston, C. McAdam, G. Mckinlay and M. Wilson drove the development of the hospital evacuations. C. Spalding worked on the design and implementation of the epidemiological models.

12. REFERENCES

- [1] Batty, M., Desyllas, J., and Duxbury, E., Safety in Numbers? Modeling Crowds and Designing Control for the Notting Hill Carnival. *Urban Studies*, (40)8:1573-1590, 2003.
- [2] Bryan, J.L., Human behavior in the MGM Grand Hotel fire. *Fire Journal*. 76:37-48, 1982.
- [3] Cocanour, C., Allen, S., Mazabob, J., Sparkes, J. Fischer, C., Romas, J. and Lally, K., Lessons Learned from the Evacuation of an Urban Teaching Hospital. *Archives of Surgery*, 137, pp. 1141-1144, 2002.
- [4] Drager, K.H., Lovas, G., Wiklund J., Soma, H., Duong, D., Violas, A. and Laneres, V., EVACSIM - A comprehensive evacuation simulation tool. In *Proc. of the 1992 Emergency Management and Engineering Conf., Soc. for Computer Simulation*, Orlando, Florida, 101-108, 1992.
- [5] Ennis, G. and Lindsay, M., VRML Possibilities: The Evolution of the Glasgow Model, *IEEE MultiMedia*, (7)2:48-51, April 2000.
- [6] Gwynne, S., Galea, E.R., Parke, J. and Hickson, J., The Collection of Pre-Evacuation Times from Evacuation Trials Involving a Hospital Outpatient Area and a University Library Facility. In D. Evans (ed.) *Proceedings of 7th IAFS Symposium WPI, USA*, pp877-888, 2003.
- [7] Jafari, M. Bakhadyrov, I. and Maher, A., Technological Advances in Evacuation Planning and Emergency Management: Current State of the Art, Technical Report to the Federal Highway Administration, EVAC-RU4474, Center for Advanced Infrastructure & Transportation, Rutgers University, New Jersey, USA, 2003.
- [8] Johnson, C.W., *A Handbook of Accident and Incident Reporting*, Glasgow University Press, Glasgow, 2003. <http://www.dcs.gla.ac.uk/~johnson/book>, 2003.
- [9] Johnson, C.W., Applying the Lessons of the Attack on the World Trade Center, 11th September 2001, to the Design and Use of Interactive Evacuation Simulations, In *Proceedings of ACM CHI 2005*, ACM Press, New York, USA, 651-660, 2005.
- [10] Johnson, C.W. and Hancock, C.P., Thinking the Unthinkable: Exposing the Vulnerabilities in the NHS Response to Coordinated Terrorist Actions, *Health Care Risk Report*, (12)5, 2006.
- [11] Johnson, C.W. Using Evacuation Simulations to Ensure the Safety and Security of the 2012 Olympic Venues. In P. Swuste and A. Hale (eds.) *Proceedings of the 3rd International Conference Working on Safety*, European Agency for Safety and Health at Work, Delft, Netherlands, 2006.
- [12] N. Latman, TCPPP Personality Profile, In *The Fourth Triennial International Fire and Cabin Safety Research Conference*, 15-18 November 2004, Parque das Nações Conference Centre, Lisbon, Portugal, 2004.
- [13] McCarthy, P.M., and Gaucher, K.A., Fire in the Operating Room—Developing a Fire Safety Plan, *AORN Journal (Journal of the Association of Perioperative Registered Nurses)*, 79:3, 2004.
- [14] Still, K., 2000, *Crowd Dynamics*, PhD Thesis, Department of Mathematics, University of Warwick, UK, August 2000.

- [15] Thompson, P.A. and Marchant, E.W., Computer and fluid modeling of evacuation. *Safety Science*, (18)277-289, 1995.
- [16] Tong, D. and Canter, D., The Decision to Evacuate: A study of Motivations which Contribute to Evacuation in the Event of Fire. *Fire Safety Journal* 9:257-265. 1985.
- [17] US National Fire Protection Association (NFPA), Summary Fire Investigation Report---Hospital Fire, Petersburg, Virginia 5 People Killed, 12/31/1994, Department Secretary for Fire Investigation, 2005.