Extending the Use of Evacuation Simulators to Support Counter Terrorism: Using Models of Human Behaviour to Coordinate Emergency Responses to Improvised Explosive Devices

Chris. W. Johnson, DPhil and Louisa Nilsen-Nygaard;
Dept of Computing Science, Univ. of Glasgow, Scotland, UK.

Keywords: public safety; evacuation; security; counter terrorism

Abstract

Software simulations have been widely used to model evacuations from fire but very few have been used to analyse a wider range of hazards, including terrorist attacks. The following pages describe how one group of evacuation simulations has been extended to support the risk assessments that drive counter terrorism. Two key areas are discussed; changes in the human response to the detonation of improvised explosive devices and the modelling of coordinated attack scenarios. These two issues create significant differences between the use of simulators to model 'conventional' evacuations and those that can be used to mitigate the effects of terrorist attacks.

Introduction

In the aftermath of the attacks in 2001, counter terrorism agencies were forced to consider an increased range of coordinated threats to public safety. The bombings in Bali, in Madrid and in London also increased awareness of the risks posed by terrorist actions. Many government organisations, especially in the US and in the UK, recruited large numbers of additional staff. This expansion was the direct consequence of political initiatives to identify and stop potential terrorist cells before there were any further attacks. However, it is widely recognised that these initiatives can never guarantee public safety. There will always remain the risk that a determined group or individual will evade detection. It is, therefore, critical that we develop a range of tools that can be used to mitigate the consequences of any potential terrorist attacks.

Computer aided design tools provide essential support across many different industries ranging from the construction of large public buildings through to the engineering of automobiles and aircraft. These software systems can help to predict the behaviour of components and materials in a variety of adverse events. For example, simulators have been used to assess crash impact damage on automobiles before prototypes have ever been constructed. Similarly, architectural models can be subjected to a range of stressors to determine the impact of earthquakes. Computational Fluid Dynamics (CFD) toolkits have also been developed to model the progress of smoke and fire early in the development of public buildings. The following pages describe initial work to extend the application of one class of simulation tools to support risk assessment in counter terrorism. The intention is to model the interaction between building occupants and public spaces in the events before and after the detonation of an improvised explosive device.

Evacuation Simulators

The computerised simulation of emergencies has been an area of research for many years, often with the aim of helping architects to avoid creating dangerous bottlenecks in building which could cost lives during evacuations. (Galea, 2006, Johnson, 2008). These tools provide important benefits. For instance, the image on the left of Figure 1 illustrates the interface to a system that was developed to model egress from a Cardiology Ward in a large hospital. The ethical problems of involving patients and staff in a 'real' exercise combined with the potential impact on continued healthcare together dissuaded staff from rehearsing the evacuation of this unit. Instead, timings were taken to assess how long it would take staff to move individual patients. This data was then used to calibrate a computer simulation of an evacuation across the entire ward. The results showed that it could take up to an hour just to move patients behind a firewall, without moving them out of the building. It also enabled hospital managers to assess the impact that reduced, night staffing levels had upon average egress times. However, these simulation tools have, typically, designed to consider scenarios that centre on localised fires or partial building collapses. They cannot easily be applied to support the risk assessments that are increasingly important within the field of counter terrorism. For example, it can be difficult to model the impact of coordinated attacks on different areas within the same target. In particular, 'conventional' evacuation simulators make strong assumptions about the availability of emergency crews that may not be sustained in counter terrorism scenarios when they themselves often form a target for secondary attacks. Similarly, significant differences have been observed in crowd behaviours in the aftermath of fires and, for instance, following the detonation of Improvised Explosive Devices (IEDs). The following sections argue that a number of further differences complicate the use of existing evacuation simulators for counter terrorism applications. The closing paragraphs explain how we have used these insights to develop an initial simulator that can be used to account for the impact of blast and fragmentation from an IED on crowd behaviours within a major public building. Although the focus of this work is on preparing for the detonation of a conventional device, the closing sections discuss on-going research to further extend our tools to consider the potential threat from Chemical, Biological, Nuclear and Radiological (CBRN) devices.

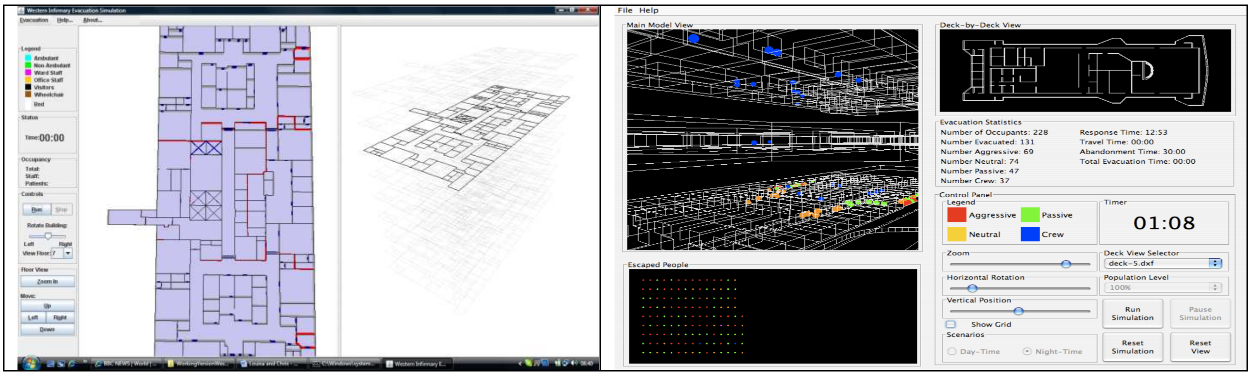

Figure 1: Examples of Evacuation Simulators (Johnson, 2006).

<u>Who Uses The Simulator?</u>

The right-hand image in Figure 1 provides a further example of an evacuation simulator that is used to illustrate the remainder of this paper. The screenshot shows the Passenger Ship Evacuation Simulator (PaSES) system. This models the evacuation of a medium sized roll-on roll-off passenger and freight vessel. The vessel illustrated in Figure 1, carries up to 190 passengers and 40 crewmembers. The software was developed in response to a number of initiatives within the maritime industries, which have realised the potential of computer simulation in helping to focus design support for potential evacuations. The International Maritime Organisation (IMO) published initial guidelines for the egress analysis of new ship designs in July 1999. The original recommendations advocated the use of a thorough hand-executed analysis before construction. This type of analysis is time consuming. It can also be error prone with repeated calculations being performed for each of the occupants within a vessel. These problems led many people to consider the use of computer modelling instead. This allows for in depth, accurate and flexible analysis to be carried out as the simulators can be easily updated and re-run to incorporate design changes on the fly (IMO, 2002).

The IMO guidelines and supporting documentation make it clear that the use of simulators can help many different stakeholders across the maritime industries including designers, operators and owners, emergency service personnel, registration and certification bodies, insurance companies etc. However, it is unclear who might benefit from the use of evacuation simulators in planning for potential terrorist attacks. While it is generally accepted that there is a need to prepare for evacuation in the case of fire, it is far less clear that the owners and operators of a vessel need to devote the same degree of care in their preparation for a terrorist attack given that such events have a much lower likelihood. This is not simply a problem for counter terrorism agencies in the maritime industry. The UK National Counter Terrorism Security Office created operation Argus to persuade retailers that they might be a future target even though they had previously been a target for the IRA. Many participants expressed a degree of fatalism over the amount that they could do to protect themselves and their customers from attack. As we shall see, simulation tools can help stakeholders to visualise what might happen during an attack on their premises. This, in turn, can help to drive future planning for what can be relatively unlikely but extremely serious incidents.

<u>What Scenarios Can We Simulate?</u>

The IMO describe two primary scenarios that should be considered during the analysis of potential evacuations. The first scenario focuses on day-time occupant distributions. The scenario involves all passengers in public spaces while two-thirds of the crew start the evacuation in work areas and one-third of the crew in public spaces. The second scenario focuses on a full evacuation from a night-time occupant distribution. All of the passengers are assumed to be in their cabins. Two-thirds of the crew should be in their cabins. One-third of the crew start the evacuation from their work areas. A range of additional constraints can be introduced to elaborate these initial scenarios. For example, maritime simulators can be used to model successively more serious damage to the vessel. They can also be used to assess the impact of different sea conditions on the course of an evacuation as passengers and crew struggle to navigate through the vessel to their assembly points.

Counter terrorism applications raise a host of questions about the types of scenarios that should be simulated. It is not simply a matter of introducing models of blast and fragmentation into an existing evacuation simulator. Most previous systems have focused on localised fires or partial structural collapses. This is justified because these scenarios account for the majority of adverse events involving large public structures. However, recent events in Iraq and Afghanistan, as well as the attacks on London and Madrid, highlight the need for a broader perspective. The principle reason for this is that most simulators cannot easily be used to consider evacuations under coordinated terrorist attacks. The threats posed by these incidents are illustrated by a suicide bombing attack on Mustansiriyah University in Iraq during January 2007. A car bomb was detonated at one of the two principle entrances to the site. This led to a partial evacuation that drew crowds to the other exit where a suicide bomber detonated their device. This is not an isolated incident. Hours before, another coordinated bombing took place in a second hand motorcycle market in the Shia Bab al-Sheik neighbourhood of Baghdad. The first blast drew onlookers and the emergency services, who were then hit by a second explosion moments later. Not only are these coordinated scenarios difficult to simulate with many ‘conventional’ software tools, they also illustrate an important reason for the use of simulators to help anticipate potential terrorist attacks. In these incidents, the use of ‘standard’ evacuation procedures that were designed to protect the public from localised fires created opportunities or vulnerabilities that were exploited by the terrorists. Building occupants, spectators and the emergency services gathered at common assembly points that were the target for secondary devices. Software simulation tools can be developed, for example, to identify ways in which evacuations can be synchronised between the decks of a vessel or between different areas of a building to help disperse the crowds that otherwise create significant opportunities for terrorist attacks.

What Results Can We Obtain?

Most of the simulators that have been developed according to the IMO guidelines structure their analytical support around three distinct phases in the response to any incident. Awareness time (A); this is the time taken for the passengers to become aware of the alarm and react. This combines a number of secondary stages including the processes involved in perceiving the alarm and the decision making processes that drive any immediate response. The latter requires that the individual is provided with the appropriate information by which to form their understanding, recognising that what they need is dependent on their existing knowledge. The model also accounts for travel time (T); this is the time taken for passengers to reach the muster points for evacuation to the lifeboats.Embarkation (E) and launch (L) time; this is the time taken to load passengers onto the lifeboats and launch them. The total time from the sounding of the alarm to the last person leaving the ship can be derived using the following formula:

$$Evacuation\_Time = A + T + 2/3(E+L). \qquad (1)$$

Only two-thirds of the embarkation and launch time is used in the final calculation as there is assumed to be a degree of overlap between each phase. For the chosen vessel, this yielded mean values of around 35 minutes for the night time scenario and 30 minutes for the day time scenario. Both were well within the allowed 80 minutes for a ship of her size. However, the results could be significantly influenced by introducing adverse weather or smoke-filled corridors into the PaSES scenarios. Successive runs of the simulator can be used to assess the impact that these and other factors, including the passenger to crew ratio, can have upon the total predicted evacuation times for a particular vessel configuration. This focus on evacuation time is justified because the aim of such an exercise is, typically, to get the passengers and crew away from the source of any fire as quickly as possible. These ‘conventional’ simulators can also be used to calculate the cumulative risk for emergency personnel as they are deployed to help occupants evacuate a public building or other complex structure. However, this is usually a secondary focus with most attention being focused on the mean time to egress across the occupant population.

The focus on 'mean evacuation times' over 'cumulative risk for first responders' is, typically, reversed in counter terrorism scenarios. One of the most salient lessons from the attacks on the World Trade Centre is that additional emergency personnel should only be deployed once the risk to their own safety has been carefully assessed. In conventional simulations, additional emergency personnel can be deployed to reduce evacuation times. However, the risk of secondary devices raises considerable questions about the exposure of police, ambulance and fire crews in the aftermath of an initial attack. Similarly, few existing evacuation simulators model the difficult triage decisions that are essential for medical personnel rendering assistance in the 'golden hour' after an incident. In summary, whilst most conventional systems focus on mean evacuation times for building occupants it seems clear that counter terrorism applications should also consider the risk exposure both to remaining occupants and to the emergency personnel who are deployed after an attack. Although 'conventional' simulators can account for this risk exposure, significant changes must be made in counter terrorism applications to consider the impact of deliberate attempts to target first responders.

Should We Develop Low or High Fidelity Simulators?

Over the last decade, two radically different approaches have emerged to the development of simulation software. The US Government's National Infrastructure Simulation and Analysis Center (NISAC), Sandia, have pioneered highly sophisticated models for the impact of terrorist attacks not just on individual buildings but on a wide range of infrastructures including electricity and food distribution. The sophistication of these models, typically, implies that they require significant computational support. They often exploit highly parallel architectures. In contrast, the images in Figure 1 illustrate a less ambitious approach. The intention is to develop low cost tools that can be rapidly reconfigured to a range of different environments and contexts. They are intended to run on conventional processors; the PC's and Macs that are available to many different stakeholders. In consequence, PaSES uses Java3D to provide a representation of the environment that is being simulated. Users can manipulate their view of the model to focus on particular occupants in different areas of the vessel. The representation is relatively simple. A wire frame model shows the principle features of the ship. The passengers and crew are represented by coloured spheres. This 'low fidelity' approach is justified because the intention is not to provide a lifelike recreation of evacuation from a localised fire or structural failure. The intention is for the simulations to support table top exercises and other forms analysis. The models help to provoke discussion amongst different stakeholders including owners and operators, emergency service personnel etc.

This 'low-fidelity, low-cost' approach to the simulation of 'conventional' scenarios including localised fires and structural collapses may not be so appropriate for many counter terrorism applications. More advanced rendering techniques including texture mapping can be used to provide a more accurate 'look and feel' to the environments in which an attack could be staged. However, as might be expected, these techniques typically incur a computational overhead that may prevent the simulations from being used in real-time on the stakeholders own machines. At a more fine-grained level, most existing simulators focus on high-level structural components of public buildings. They accurately model the position of principle architectural features including walls, doors, ceilings etc. However, they often do not include temporary structures, including the concrete and metal barriers that have been recently deployed across many airports, railway stations etc to deter terrorist attacks. Similarly, 'conventional' simulations do not account for the non-permanent structures including kiosks, advertising panels and public art works that can either be fragmented by an IED or may help to absorb the impact of any blast.

What is the Scope or Environment of a Simulation?

There can be considerable benefits when evacuation simulators are tightly integrated with existing design processes. This enables changes to be made early in the development cycle while the costs incurred by any changes are relatively low. It is for this reason that the PaSSES system is based around the same dxf file format that is used in leading CAD/CAM packages. In other words, the output of design tools can be used to drive the simulations without the need to develop costly specialist 3D models. Once loaded into the application a logical representation is derived from the graphical model. A graph of linked vertices is created; each one represents a point within the environment which can be occupied by passengers and crew. The graph also supports collision detection; simulated occupants should not typically be allowed to walk through solid walls. Most simulators, like those shown in Figure 1, focus on single structures such as individual buildings or vessels. It makes little sense to extend a model beyond the confines of the structure that a designer or architect is currently working on. They may only have limited opportunities to affect the architecture of surrounding structures in their environment.

This focus on individual buildings makes sense when the companies and organisations that pay for evacuation simulators are principally concerned to protect the safety of the occupants in the buildings and other structures that they are paid to construct. However, recent terrorist attacks have shown that this is far too narrow a view for most counter terrorism applications. The 2002 Bali bombings involved three devices that affected a network of streets in the same district. These included a backpack device carried by a suicide bomber and a large car bomb that were both detonated close to nightclubs in Kuta. The 2005 Bai bombings repeated this pattern. One of the three bombs was detonated in a crowded in the main square at central Kuta. Two more went off on Jimbaran beach. Three further unexploded devices were found in the same area and apparently failed to go off after the security forces shut down the island's mobile telephone network following the initial blasts. Similarly, the attacks on the World Trade Centre had knock-on effects that went well beyond any single structure. The aircraft impacts undermined building seven as well as the Twin Towers. This was a 47 stories tall structure that was connected to the World Trade Center plaza by an elevated walkway. Building seven was damaged by debris and by fires which burned throughout the afternoon leading to its eventual collapse. The recent suicide attacks in Baghdad exploited detailed local knowledge not just of individual buildings but also of the ways in which people move between adjacent areas. As we have seen in previous attacks, the scale and planning of potential incidents has increased so that multiple coordinated attacks may take place at the same time across a city or district. In such circumstances, simulators must be extended beyond the walls of a single building or vehicle. The closing sections of this report will describe a new generation of simulators that focus on a higher granularity of environmental model. These systems can be used to analyse the impact of coordinated attacks on urban and national infrastructures. They have also been extended to consider the consequences of knock-on failures, for example in the electrical infrastructures, which now extend across national borders. An important difference with the high-fidelity approaches mentioned in previous sections is that these simulators are also intended to run on the computational resources available to most stakeholders.

What Aspects of an Emergency Response Should We Simulate?

Most existing evacuation simulators focus on the egress of occupants under a range of different scenarios. As mentioned previously, the aim is typically to identify ways of reducing the mean times to get as many people as possible out of a structure. In other areas, the requirements are more specific. For example, aviation simulators must typically determine whether a particular aircraft meets FAA/JAA 25.803; 'For a airplanes having a seating capacity of more than 44 passengers, it must be shown that the maximum seating capacity, including the number of crew members required by the operating rules for which certification is requested, can be evacuated from the airplane to the ground under simulated emergency conditions within 90 seconds'. There are further exceptions. For example, many hospitals use a form of 'horizontal evacuation' where patients are not moved outside a ward but are instead moved behind fire walls until the emergency services arrive. However, many hospital and aviation simulators tend to focus on the movement of building occupants rather than on the intervention of fire-fighters or other rescue services.

An important requirement is that simulators consider the intervention of emergency personnel in counter terrorism applications. The types of risk assessment that must be conducted before moving onto a particular site are very different in counter terrorism scenarios compared to fires and other forms of partial structural collapse. The hazards of secondary devices must be considered, especially when there is a risk that emergency personnel might be deliberately targeted. The usual priorities for egress may not be appropriate; 'invacuation' techniques may be used to move individuals inside any structure that can provide protection from further blasts or debris. These additional hazards can be compounded by the different chains of command and coordination that are necessary for any response to such incidents. The focus of the rescue services on assisting the injured must, typically, be considered alongside the requirements of other agencies to preserve the crime scene and apprehend any remaining terrorists. Above all, there is a need for dynamic simulators to help identify potential vulnerabilities in evacuation plans. For instance, horizontal evacuation techniques offer an appropriate response to 'conventional' incidents. Instead of moving occupants as far away from a hazard as possible, they are left behind fire walls until emergency personnel arrive. This avoids staff having to move patients between floors in many hospitals. However, this approach introduces significant vulnerabilities if the emergency response is delayed, for example, by the detonation of secondary devices (Johnson and Hancock, 2006).

What Aspects of Occupant Should We Simulate?

The developers of evacuation software must simulate the problem solving processes that individuals and groups use to navigate towards particular exit points. If the occupants are already familiar with the layout of a location then this can be relatively straightforward, although some account must be made for the impact of smoke and partial structural collapse. However, the simulation of the cognitive processes involved in navigation can be considerably more complex if the occupants have to traverse unfamiliar environments. AI techniques can be used to simulate the problem solving techniques that individuals apply in these situations. For example, PaSES uses the A* path-finding algorithm. This is an approximation. It cannot accurately model the range of behaviours that are often observed during complex evacuations. For example, previous studies have shown that many individuals will choose to go out the way they entered a building even though this may lead them to walk passed fire exits that might provide a more rapid egress (Johnson, 2006). This is often justified by the occupant's uncertainty over such exit routes; they are worried that they might be led back into a fire or that the doors might be locked and so on. A further limitation with the use of these relatively simply navigation algorithms is that they often ignore the impact of crowd behaviours on individual decision making. For instance, 'flocking' can persuade individuals to follow larger groups rather than move in the direction that they might otherwise take. The impact of such phenomena remain a considerable topic for debate, hence, it is difficult to determine the extent to which they might be relied upon in evacuation simulations.

Navigation is only one aspect of the complex behaviours that must be considered even within 'conventional' evacuation simulators. The PaSES system uses Monte Carlo techniques. An initial questionnaire is issued to identify the physiological, perceptual and cognitive attributes of building occupants. This determines the percentage of people who are physically fit. It can also be used to assess the proportion of that are more aggressive compared to the other building occupants. When an individual is added to the model then a random number is generated. Supposing that 50% of the population were found to be 'aggressive', if the number were in the range from 1 to 50 then that individual would be classified as aggressive. If 20% were found to be 'neutral' and the random number was between 51 and 70 then they would be classified as neutral and so on. This process ensures that the composition of the population varies as the random numbers change between each run of the simulation. Over time, however, the distributions will tend towards those that were identified from the initial questionnaire. The same techniques can be extended to include a range of different cognitive characteristics that are assumed to have some impact on evacuation behaviours.

Monte Carlo techniques are not only used to determine the cognitive, perceptual and physiological attributes of the building population, they are also used to approximate the behaviour of those individuals during an evacuation. Probabilities are associated with the likelihood of individuals taking particular actions within the next interval of time. The probabilities of individual actions are conditioned by the attributes that were assigned during the population phase, described in the previous paragraph. For example, the developer of a simulation might specify that there is an 80% probability that an 'aggressive' individual will push forward if they have the opportunity to move towards an exit. If a random number falls between 1 and 80 then they will advance. If the number falls outside this range then they will not move. Another individual may then move into the free space providing that the random number associated with that action falls within the bound of their associated probability of movement. This approach makes it more likely that 'aggressive' individuals will move first, although this is not always the case.

More complex models of human behaviour have been developed. However, the wider the range of cognitive characteristics that are represented in a model then the harder it can be to validate any associated behaviours. In other words, if we introduce parameters such as 'assertiveness', 'risk aversion' and 'fear' then we must explain how these attributes influence an individual's behaviour during any evacuation. We lack empirical studies that can be used to accurately predict how such parameters might impact individual actions under highly stressful situations. This lack of validation is even more problematic for counter terrorism scenarios. CCTV footage of previous attacks can offer significant insights into individual and group behaviours in. However, these images provide limited opportunities to identify the psychological and physiological markers that might help us predict likely behaviours in the aftermath of future incidents involving different groups of people in different locations. There are significant differences in the ways in which crowds reacted after the Bali nightclub bombings compared to those that were caught up in the blasts on commuter trains in Madrid. We might anticipate further differences between these attacks and, for instance, any future incident at sporting venues or shopping malls where family groups might be involved. Some evidence can be gleaned, for example from crowd behaviours during previous stadium disasters such as the Bradford fire. These approximations and analogies can be dangerous allies. For example, the gradual development of a fire can prevent many in the crowd at sporting venues from realising that they are in danger. However, the noise and blast associated with an IED can have an immediate impact on individuals who are a long way from the source of any detonation.

What Aspects of Physiology and Injuries Should We Simulate?

Conventional evacuation simulators often account for the physiological differences that characterise the populations in a particular environment. Previous sections have explained how these systems can be used to model the egress of groups that are, typically, excluded from participation in 'live' evacuation drills. Models can be calibrated to include family groups containing the very young, the elderly or patients in healthcare settings. The inclusion of these individuals is an important strength of evacuation simulators because the consequent physiological differences can have a profound impact on evacuation times. They can also increase the likelihood of bottlenecks as parents pause to pick up a child or as other members of a crowd move past a slower occupant as they move towards an exit. Evacuation software can capture not only the increased likelihood of crush injuries to particular population groups, it can also be used to represent the increased consequences of those injuries over time. For example, crush injuries will typically lead to shock, hyperkalaemia which may precipitate cardiac arrest, hypocalcaemia etc. These all tend to have a greater impact on the young and the elderly increasing the need to focus medical assistance on their treatment in the aftermath of any evacuation (Greaves, 2004).

Simulators can also account for the increased susceptibility of particular individuals to smoke inhalation during an evacuation. Most fatalities during fires stem from irritants, including hydrochloric acid, sulfur dioxide, oxides of nitrogen and ammonia or from asphyxiants (toxicants), including carbon dioxide, hydrogen cyanide, carbon monoxide and hydrogen sulphide rather than from burns. These different agents seem to have a higher impact on the elderly and on the young; the effects are faster and more serious than on other groups. Again, however, there is a lack of validation for many of the models that can be used to differentiate between the impacts of these different factors across populations. There is no agreed repository for data about morbidity in fires or other forms of evacuation. Studies, therefore, have to extrapolate from local samples that contain numerous forms of bias that make it difficult to support comparisons between similar work from different countries. Further problems arise when individuals are affected by smoke but they still succeed in escaping from the building, in such instances there may be no longer term record of the impact of inhalation either on them or on the other building occupants.

Smoke inhalation and crush injuries are only part of a wider spectrum of injuries that must be considered when simulators model counter-terrorism scenarios. It is important to develop appropriate models for the blast and fragmentation injuries that may result from terrorist devices. This is a difficult and controversial task (US DARPA-NRC, 2004). There are considerable disagreements over the power of IEDs from the models that have been developed by military and civil agencies; this is discussed in greater detail in the following section. Even if we can reach some consensus on the likely force generated by an IED, it is again difficult to validate the models that can be used to analyse the impact that this force will have upon building occupants. It is important to consider both the direct blast but also the effect of any fragmentation devices that are packed around the explosive. Crowds are also often injured by the secondary effects of fragmentation caused by the impact of an initial blast on any surrounding fixtures and fittings. These problems are exacerbated by the use of multiple devices. Modelling is also complicated by the dissipation of force and fragmentation when there are other objects between the blast and a target. In spite of these difficulties, a number of models have been developed that can relatively easily be introduced into computational simulations (Kress, 2005). Elementary geometry can be combined with stochastic estimates of injuries for different crowd densities. Such models confirm the empirical observations that increasing crowd densities may reduce the fatality rate for some types of IED. If a dense crowd is present around a suicide bomber then the number of expected casualties may actually decrease due to the effect of crowd blocking. Those closer to the bomber will stop the majority of the fragments from reaching those further away, thereby reducing the number of serious injuries and fatalities.

Can We Predict Blast and Fragmentation?

The previous section argued that one of the principle differences between convention evacuation simulators and those that can be extended to model counter terrorism scenarios is that the latter must account for the impact of blast and fragmentation on bystanders and on surrounding structures. This is a difficult, if not impossible, task given the enormous variations between improvised, commercial and military grade explosives. Further complexity arises because their chemical composition changes over time in response to changes in supply but also through technical innovation between terrorist groups. The modelling of blast and fragmentation is exacerbated by the use of multiple devices and by considerable differences in the quantity of explosives that can be used. As we have seen, previous attacks have used combinations of suicide bombers, car bombs and devices hidden in lockers or other semi-permanent structures.

The modelling of explosives within any simulator also relies upon assumptions about the composition of any explosive device. For example, many of the IED used in Iraq are simple platter charges. These are constructed from several kilograms of plastic explosive pressed into a similar mass of flat metal, typically steel. This will propel the platter into a target with an approximate velocity of 1,800 m/s at up to 50m. For other targets, Explosively Formed Penetrators (EFPs) have been deployed. In these devices, the force of a blast helps to form a penetrating projectile that can be effective more than 80m from the target. For example, cylindrical shaped charges can be tipped with a concave metal disc, typically made of copper. Variants on this type of device have been successfully deployed against Abrams M1A2 tanks. US Army field manual FM20-32 provide a useful starting point for the development of counter terrorism simulators because it provides an initial taxonomy for improvised explosive devices (US Army, 1998). It distinguishes between high-explosive, artillery-shell antitank devices, platter charges, improvised Claymores, grapeshot antipersonnel devices and barbwire antipersonnel devices. Although the focus is devices that have been used against organised military units, all have been used on civil populations in different parts of the globe. We are also constructing a more focussed taxonomy that extracts key features of IED incidents reported by media and intelligence services to help counter terrorism agencies identify common features between different attacks over time.

Can We Identify Countermeasures?

The modelling and simulation of events during and after terrorist attacks is very different to simulating other emergency situations such as fires. The modelling of the progression of fire and its by-products through a structure is a relatively mature area even though many problems remain, for example, in the finite element analysis of smoke. Extensive research has been undertaken into designing buildings to manage the effects of fire, protect occupants, and help emergency workers to render aid. Techniques such as positive pressure ventilation, which allows fire fighters to control the flow of smoke through a building by means of doors and windows, have been aided by simulation software. In some places simulations have even been used to determine which windows emergency workers can safely break to gain access to the building without causing further spread of the fire. The key benefit of computer simulations in this domain is that it is possible to anticipate the interaction between building occupants and their environment as they face different types of hazard. In the past, it was difficult to plan for the deployment of these dynamic counter measures as it was difficult to predict which areas should be ventilated as occupants moved towards the exit. In contrast, software tools have been used to show that particular countermeasure strategies provide adequate protection under a wide range of scenarios.

Software tools can help to design buildings that protect their occupants from possible terrorist attacks. A number of specialist packages assess the impact that blast can have upon different structures. The same tools can be used to identify potential benefits from additional reinforcement, reducing the amount of glass frontage etc. Existing CAD/CAM tools have also been used to plan the vehicle and occupant flow measures that are increasingly being used to control access to possible targets. These can be seen at a local level in the tons of concrete that have been used to prevent vehicles from getting close to UK airports following the attack in Glasgow during June 2007. Similarly, a ten metre stretch of pavement outside the Scottish Parliament building is now protected by over 60 concrete bollards. Unfortunately, the risk assessments that can be informed by these tools must be continually revised as the resources available to a potential terrorists change over time. Similarly, they may also be revised in the aftermath of successful attacks in other areas of the globe. It is for this reason that the current forest of bollards is being replaced outside the Holyrood Parliament building.

The previous paragraphs have described the use of ‘static’ models for counter terrorism applications. In other words, they assess the impact of permanent and semi-permanent changes in the built environment. They help to plan the use of concrete structures to restrict access or to locate blast walls that limit the consequences of any attack. One limitation of these static models is that the introduction of counter terrorism measures can create other vulnerabilities. For example, the deployment of concrete structures can prevent the emergency services from getting close to public buildings in the aftermath of an attack or during more ‘conventional’ incidents. The introduction of blast walls and of other semi-permanent structures can hinder the evacuation of building occupants. Unfortunately, there have been relatively few attempts to develop more 'dynamic' models of the interaction between occupants and the buildings under attack. In consequence, there is the strong possibility that some of the changes we have introduced into public buildings will actually create opportunities for future terrorist attacks. The bottlenecks at airport security barriers could be exploited by a secondary attack. The creation of 'sterile forecourts' in front of airport buildings can add to crowds around check-in desks. These vulnerabilities can be systematically analysed using dynamic simulators that model the movement of people within these environments during normal conditions and in the aftermath of an initial attack.

The Railway Station Case Study

Central Station is the busiest railway station in the UK outside of London, with an annual throughput of over 34 million people. It has a peak weekday occupancy of more than 15,000 people. It consists of two parts; the High Level station from which inter-city trains to the South and some local trains run, and the Low Level station through which local services run. The High Level station also houses several shops, a pub and a hotel. As the busiest transport hub in Glasgow, it is considered as a potential terrorist target. Each year a table top exercise is held. This involves around 60 staff from the station, transport police and the train operating companies. The exercise is designed to prepare for possible attacks and to help refine the procedures in place for dealing with them. We have, therefore, used the findings of the analysis presented in this paper to inform the development of a dynamic simulation tool for counter terrorism scenarios. The intention is that this tool can be used by staff to support these annual exercises, for instance, by working through the inter-agency response to a range of different scenarios. A number of initial incidents have been identified. These include:

1. Suicide bomber(s) on the main concourse;
2. Suspect vehicle entering through main/side entrance;
3. Suspect vehicle entering through car park entrance/exit;
4. Coordinated use of 2 or 3 (described above) followed by attack scenario 1.

Key issues to be addressed include the likely numbers of casualties for each scenario at different times of the week, the exits most likely to be used by passengers and any subsequent bottlenecks that might be vulnerable to secondary attacks, etc. The simulation work is focussing on three initial stages; the first is the normal operation of the station with passengers coming and going through the station; the second is the detonation phase when initial casualties can be seen; and third is the evacuation phase where the reaction of the passengers to the first explosion can be seen.

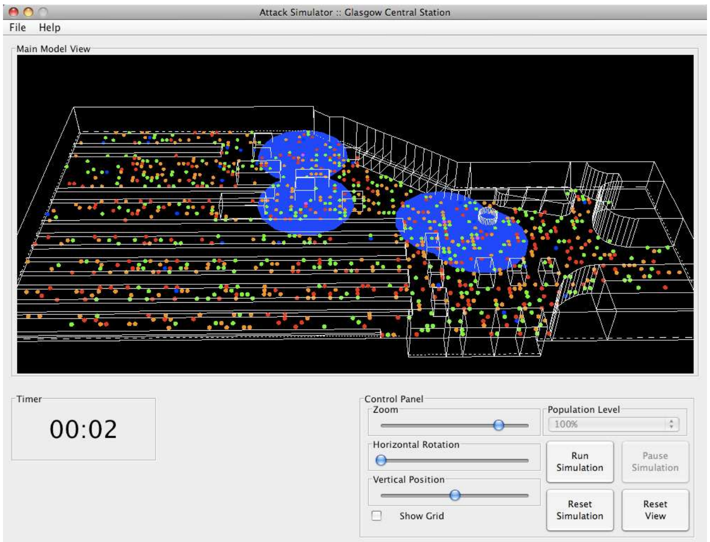

Figure 2: Interface to an IED Simulation

Figure 2 illustrates the interface to the prototype application; passengers are distributed throughout the station. In this instance, suicide bombers can be identified by the circles that represents the potential targets that could be caught in any blast. This changes for each bomber as they and the other passengers move throughout the station concourse in real-time. As the bomber moves around the station the other agents will enter and leave these 'danger zones'. The size of the blast and fragmentation areas can be varied to allow for larger and smaller devices given the type of explosive used. By monitoring the number of agents inside the zone at any one time and using the formulae in Kress (2005) we can continually update the estimated number of casualties. As mentioned, these equations account for the mitigation of blast by objects between the source of any explosion and a potential victim. In the future, it is hoped that this will provide important information for the security services that *in extremis* may be forced to act against a suspected bomber as they move through public spaces. The simulator does not account for concussive force of the blast on the people and structures in the surrounding area, although this is an obvious area for further development.

An important consideration in the simulation is the positioning and control of the agent representing the suicide bomber. If the bomber were stationary and randomly placed, this would eliminate both the ability to monitor changing casualty rates and the element of interactivity which could make for a useful training tool. A better option would be to have the bomber act as part of the crowd of passengers moving through the station, or for the user to be able to nominate any member of the crowd as the bomber by clicking on them in the 3D view. A third option is for the user to directly control the movements of the bomber inside the simulation. Further work is required to determine which of these approaches might yield the greatest benefits to different stakeholders.

Conclusion and Further Work

There are many situations where it can be difficult to conduct the drills and exercises that help to prepare for adverse events including fires and structural collapse. For example, there are clear risks in rehearsing the evacuation of hospitals and care homes when exercises may place patients' lives at risk and interrupt the provision of medical services. Similarly, it can be hard to conduct evacuation exercises, for example in underground transit systems or in shopping malls, when they involve young children or the elderly. Software simulations can be used to address these problems. They can be calibrated to model a broad cross section of the population including the elderly and the very young that are often excluded from evacuation exercises for ethical reasons. These tools have been widely used to model evacuations from fire but very few have been used to support a wider range of hazards, including terrorist attacks. This paper has described how evacuation simulations can support counter terrorism applications. An important benefit of these tools is that they can be used to explore alternate intervention strategies as part of the training and planning of emergency service teams. However, these simulators must account for the diverse human response to the detonation of improvised explosive devices. They must also be capable of modelling of coordinated attack scenarios. These two issues create significant differences between the use of simulators to model 'conventional' evacuations and those that can be used to assess the risks posed by terrorist attacks on civilian targets.

Acknowledgement

This work was produced as part of an EPSRC/BAE Systems CASE studentship. Thanks are due to anonymous reviewers inside BAE who provided valuable feedback on an initial draft of this work.

References

E.R. Galea, Proposed Methodology for the Use of Computer Simulation to Enhance Aircraft Evacuation Certification. AIAA Journal of Aircraft, Vol 43, Number 5, pp 1405-1413, 2006.

I. Greaves, Consensus statement on crush injury and crush syndrome, Accident and Emergency Nursing, Volume 12, Issue 1, January 2004, Pages 47-52, 2004.

International Maritime Organisation, Interim Guidelines for Evacuation Analyses for New and Existing Passenger Vessels, Ref T4/4.01, Circular 1033, London, UK, 6 June 2002.

C.W. Johnson, Using Evacuation Simulations to Ensure the Safety and Security of the 2012 Olympic Venues, Safety Science (46)2:302-322, 2008.

C.W. Johnson, The Application of Computational Models for the Simulation of Large Scale Evacuations Following Infrastructure Failures and Terrorist Incidents. In D. Skanata and D.M. Byrd (eds), Proceedings of NATO Research Workshop on Computational Models of Risk to Infrastructure, 9-13 May 2006, NATO, Primosten, Croatia, 2006.

C.W. Johnson and C.P. Hancock, Thinking the Unthinkable: Exposing the Vulnerabilities in the NHS Response to Coordinated Terrorist Actions, Health Care Risk Report, (12)5, 2006.

M. Kress, The Effect of Crowd Density on the Expected Number of Casualties in a Suicide Attack, Naval Research Logistic, Vol. 52, No. 1, 22-29, 2005.

National Research Council of the National Academies (2004) Existing and Potential Standoff Explosives Detection Techniques (Natl. Acad. Press, Washington, DC).

US Army, Mine/Countermine Operations, FM 20-32, Headquarters, Department Of The Army, Washington, DC, 29 May 1998

Biography

Chris.W. Johnson, DPhil, MA, MSc, FBCS, CEng, CITP, Department of Computing Science, University of Glasgow, Glasgow, G12 8RZ, Scotland, UK, telephone +44 (141) 330 6053, facsimile +44 (141) 330 4913, e-mail – Johnson@dcs.gla.ac.uk, web page http://www.dcs.gla.ac.uk/~johnson

Chris Johnson is Professor of Computing Science at the University of Glasgow in Scotland. He heads a small research group devoted to improving the reporting and analysis of incidents and accidents across safety-critical domains ranging from healthcare, to the military to aviation and rail.

Louisa Nilsen-Nygaard, Department of Computing Science, University of Glasgow, Glasgow, G12 8RZ, Scotland, UK, telephone +44 (141) 330 2000 ext 0917, email louisa@dcs.gla.ac.uk, web page http://www.dcs.gla.ac.uk/~louisa

Louisa Nilsen-Nygaard is PhD student in the Computing Science Department at the University of Glasgow in Scotland. Louisa is doing research into using computer simulation to determine counter defences to improvised explosive devices.