

Modelling of Behaviours in Response to Terrorist Activity

D. Keep^a, I. Piper^a, A. Green^a and R. Bunder^a

^a*University of Wollongong, Northfields Ave, Wollongong, NSW 2500*
Email: ian@uow.edu.au

Abstract: In this paper we present an approach to the modelling of human interaction in complex environments and its application to a security related scenario; the evacuation of a railway station subsequent to the detonation of an improvised explosive device.

The intent of the experiments reported in this paper is to investigate the application of our existing software capabilities to the proof-of-concept scenario described above.

Our simulation framework, Simulacron, allows the development of multiple interacting modules which address matters such as motivation, scheduling and movement, controlled both by internal goals and external influences. The ability to integrate this interaction modelling with our existing, rich palette of capabilities makes this approach to simulation building of particular interest.

In this paper, we pay particular attention to the motivation module, which allows the coupling of internal characteristics with external conditions to control behaviour.

Sydney's Central Station is represented in this scenario with a directed graph of some 200 locations. Through this, five thousand commuters move over the course of two simulated hours at a temporal granularity of 5 seconds. Typical run time was on the order of one hour on a consumer desktop using a single thread.

Commuter behaviours, which have been selected based on actual, observed behaviours in similar circumstances, are represented by a number of states, driven by transition rules using a set of hypothetical characteristics.

We report the results of several variations to the base scenario where we modify commuter and assailant characteristics. Additionally, we report on a set of simulations in which simple crowd management strategies are implemented to investigate their potential impact on the evacuation process.

We conclude that the results broadly agree with expected behaviours and justify further development. This would include the incorporation of psychologically-grounded characteristics and the extension of the work to a full validation scenario.

We foreshadow a range of future initiatives including a study to determine the interaction between evasion resulting from perceived assailant threat and fatality rates.

The work to date has provided us with a simulation environment into which we can now introduce more experimentally derived parameter values, covering not only internal states but also allowing the introduction of passenger schedules and the incorporation of incoming and outgoing trains. The incorporation of additional features, such as: more realistic evasion, security force interdiction, multiple assailants including decoys, caring for the injured, families, biological imperatives, such as hunger, and disorientation during evacuations, is planned for the future.

Keywords: Simulation, evacuation, terrorism, behaviours.

1 INTRODUCTION

Microsimulation is a discrete simulation technique which allows for the modelling of the behaviour of individuals in a complex system (Connor *et al.*, 2000; Cohen, 1991; Merz, 1991). This paper builds on our previous work (Piper *et al.*, 2009, 2010; Green *et al.*, 2010) in which we describe our approach to modelling systems involving many individuals and their interactions. Our modular microsimulation framework, Simulacron, allows us to quickly build simulations with varying requirements. As described in previous publications, it has already been applied to problems involving epidemiology and terrorism.

It should be noted that the work described here involves the incorporation of new functionality into a generalised framework and is not a new, standalone package. Thus, the newly introduced functionality can be combined with our existing capabilities permitting, for example, the incorporation of psychological factors into an epidemiological simulation, greatly enhancing the framework's modelling abilities.

2 SOFTWARE: FRAMEWORK AND MODULES

Our microsimulation-based approach to modelling divides the world into entities of two classes: "cells" and "peeps". Cells are abstract locations which have no intrinsic meaning or properties whilst peeps are abstract things whose only intrinsic property is their location (a cell). Peeps are not agents; they are not autonomous but typically require far less computing resources. Cells and peeps have meaning associated with them via "fields" which are named, arbitrary chunks of data. This system allows us to extend the existing abstractions with new meaning without having to modify the simulation framework. Note that, by design, cells are isolated from one another. This means that interactions can only ever be within a cell.

The framework supports the use of "states" which allow alternate field value sets to be associated with an entity based on their current state. As an example, a peep might have two different schedules: one for use in the "working" state and one for use in the "weekend" state.

The actual simulation program is the combination of the simulation framework and one or more modules. Together, these modules enable the specific simulation behaviour. The framework provides modules with various services such as loading and management of the input data sets, reporting, messaging, etc. To date, we have modules implementing cyclic and one-off scheduling, non-deterministic dispersion, spread of infectious agents and a terrorist behaviour prototype. A user is free to use whatever subset of modules they need; the simulation need only be as complicated as the scenario demands.

The scenario chosen for this paper required the use of most of our existing functionality, which is very briefly described below. Space limitations prohibit a thorough description of the modules involved. Details of these modules can be found in Piper *et al.* (2010) and Keep *et al.* (2011); the motivation module in particular benefits from a more comprehensive description.

Event Scheduling - allows the triggering of an action on an individual peep at a specific moment in time. The module was extended to allow event schedules to be used with cells in addition to peeps.

Dispersion - this module allows for non-deterministic movement of peeps; it simulates "wandering". It also provides basic rate-limiting controls, masking and biasing.

Induced Peep Actions - this is a new module developed for this scenario. It is a generalisation of the dispersion mechanism, allowing a cell to induce any action in a peep; dispersion can only induce movement. It has the same masking and biasing features as the dispersion module.

Motivation - the motivation module allows peeps to be affected by their environment and other peeps near them. Peeps can have personal characteristics, interactions which cause changes in the characteristics of themselves and the peeps around them, and rules which drive changes in state.

Several extensions were made to the motivation module for the purposes of this scenario. The first was to allow cells to participate by giving them characteristics and interactions, allowing them to influence peeps. This was found to be necessary in order to correctly model environmental influences.

The second was to allow arithmetic expressions for interactions. Interactions are how a peep influences another peep's characteristics. Previously, interactions were restricted to a simple delta change involving, at most, two characteristics called "cause" and "effect". A static "strength" value could also be used.

As modified, interactions now specify an expression involving an arbitrary set of characteristics derived from the peep doing the influencing, the peep being influenced and the environment (the local cell). They can also use a number of derived variables including the summation, min, max or mean of a given characteristic across all peeps located in the same cell.

Infection - the infection module typically allows modelling of the spread of an infectious disease. In this scenario, we use the module to simulate the effect of an explosive device.

Transport - this module provides goal-based movement for peeps, allowing them to navigate between cells. The cells are arranged in a directed graph, implemented via an additional field on each cell involved.

3 THE SCENARIO

The scenario chosen for this paper was of a hypothetical bomb blast at Central Station in Sydney occurring during the course of two hours of simulated time, with a temporal granularity of five seconds. Normal commuter traffic is simulated for a full hour prior to the assailant's entry to the station. The assailant then proceeds to a target platform before detonating a low-yield backpack IED six minutes after entry. The simulation then continues until completion. This is intended to be a proof-of-concept to assess the feasibility of our approach and the utility of our tools.

3.1 Geometry

The scenario is limited to modelling the upper floor of the station. This includes fifteen platforms and nine ingress/egress points and associated concourse areas, three cafés, two shops, two bathrooms, ticket dispensers and windows (including associated queues), and two platform access barriers. The layout was designed to closely resemble reality with an acceptable level of detail for the simulation.¹ In total, 225 cells were used to model the layout of the station.

With the exception of platform 1, platforms are arranged in pairs (2 and 3, 4 and 5, etc.) where the platforms in a pair are connected to each other. Platforms 4 through 15 are behind the platform access barriers which restrict traffic flow between the platforms and the rest of the concourse. Platforms 8 and 9 are each further subdivided into 15 cells owing to platform 9's status as the assailant's target; a level of detail not required for the remaining platforms.

All the cells which make up the station are linked together into a directed graph; this allows commuters to employ the transport module to perform navigation around the station. The dispersion module's rate limiting features are used for the queues and barriers. Each cell also has several masked, induced peep actions which are used to model the random movement of hysterical commuters and direct certain evacuating commuters to the nearest exit. Ingress to the station is done via an event schedule on each commuter; prior to their arrival, they are kept in a holding cell which is not directly connected to the station.

The detonation of the IED causes all cells in the simulation to switch into an "aftermath" state whereby they begin to broadcast a danger level as a motivation characteristic.² This is the primary driver in subsequent peep behaviours. The danger value varies inversely with distance to the blast, being 1 at the centre of the blast and zero infinitely far away.

3.2 Peeps

In this scenario, we simulated five thousand individual commuters over the course of two hours, which represents approximately half the typical rush hour traffic.³ Each commuter was given the following randomly selected properties: an entrance/exit location, an arrival time at the station and a target platform. In addition, each commuter *may* have each of the following additional intermediate goals: one of the bathrooms, the ticket window or dispenser, a café, and a shop. Commuters were uniformly distributed among entrances and among platforms. The proportions of commuters visiting optional goals were rough estimates, based on personal observations.

¹This level of detail is similar to that used in our previous studies.

²In fact, the switch into the "aftermath" state is triggered by a pre-determined event schedule, timed to go off at the same time as the blast. This is currently the easiest way to fake a local event causing global effects.

³The number of commuters, with the given level of complexity, is presently limited by bugs in the compiler we are using; efforts are underway to circumvent these problems.

Behaviours are driven by the motivation module. There is no direct interaction between individual commuters⁴. Each commuter perceives a different level of danger, based on their awareness, the aggregate behaviour of nearby commuters and the inherent danger of the cell. As the perceived danger level rises above a minimum level, they move into an information-seeking state; this results in the commuters moving toward the blast in order to “find out what’s going on”. Once the perceived danger exceeds a critical level, the commuters then transition into one of several states (such as panic, evacuation and assistance) based on their motivation parameters. The motivation parameters and states are detailed in Section 3.3.

A single assailant enters the station one hour into the simulation at entrance 1 and proceeds to subdivision 8 of platform 9. There, he waits until one hour and six minutes into the simulation at which point he detonates, killing or injuring himself and everyone in the same subdivision and triggering the aftermath, detailed below. The blast itself is modelled as an infection: it has instantaneous transmission, no latent or asymptomatic time and is either immediately lethal to a peep or leaves said peep in the symptomatic state for however long it takes them to die.⁵

The assailant, whilst still alive, uses the motivation module to broadcast a “threat” characteristic. This may cause other peeps, based on their situational awareness, to temporarily go into an evasion state whereby they attempt to distance themselves from the assailant. Currently, this threat perception is cumulative and decays over time but requires further refinement. This evasion behaviour is extrapolated from some of the interviews quoted in Addley (2011).

3.3 Commuter Motivation

As mentioned above, commuter behaviour is driven by the motivation module described in Keep *et al.* (2011), derived from the BDI model of Georgeff *et al.* (1999). External conditions and internal characteristics, together, drive behaviour, represented by changes in each commuter’s state. The behaviours chosen for this scenario are based on our previous research, specifically the as-yet unpublished Master’s thesis by Davies (2010), itself based, in part, on the work of Drury *et al.* (2009) and Addley (2011).

The states a commuter can be in are: *<default>*, the initial state in which the commuter moves to their goal normally; *<evading>*, where the commuter is evading a threat; *<information-seeking>*, where the commuter is trying to determine what has happened; *<catatonic>*, where the commuter is stationary and unresponsive; *<hysterical>*, where the commuter moves about randomly with no fixed goal; *<assist-primary>*, where the commuter, if medically trained, moves to assist at the blast site; *<assist-secondary>*, where the commuter, having no medical training, sees at least one commuter in *<assist-primary>* and moves to assist⁶; *<evacuating-at-speed>*, where the commuter is moving toward the nearest exit; *<evacuating-normally>*, where the commuter is moving toward their original point of entry and *<injured-or-dead>*, where the commuter was caught in the blast. What states a commuter can be in and how they transition between them are determined by the commuter’s motivation parameters; except for *<injured-or-dead>*.

Each commuter has a number of characteristics, divided into two groups. *Innate* characteristics are those which have a fixed value for the duration of the simulation, whilst *derived* characteristics are those whose value is dynamically determined from the commuter’s own innate characteristics and other environmental factors. The specific states a commuter can go into depend on these. The innate characteristics determine the set of possible states, whilst the derived characteristics control which of those they are in at any given time. The innate characteristics are named: *awareness*, *empathy*, *extraversion*, *impatience*, *sang-froid* and *training*, whilst the derived characteristics are named: *danger-perception*, *leadership-perception* and *threat-perception*. It should be noted that these characteristics are plausible, hypothetical constructs invented by the authors, rather than rigorously researched psychological properties.

Rules govern which state the commuter will be in at any given moment; their form is detailed in Keep *et al.* (2011). They define actions the commuter will take when they change state, typically a change of goal. Note that the *<injured-or-dead>* state is not set by the motivation system. From the *<default>* and *<evading>* states, commuters can change into any other state. From other states, the com-

⁴The *<assist-secondary>* state is presently driven by an aggregate interaction, not a person-to-person one.

⁵Treatment is not modelled at present. The assisting commuters, detailed later, are intended to eventually fulfill this role.

⁶We acknowledge the observed behaviour whereby untrained commuters spontaneously assist the injured, but have not attempted to model this at present. Its inclusion depends on an accurate psychological model.

muter cannot return to $\langle \text{default} \rangle$ or $\langle \text{evading} \rangle$. Given the order in which the states were listed above, $\langle \text{information-seeking} \rangle$ can transition into any later state. Of the states after $\langle \text{information-seeking} \rangle$, the two $\langle \text{evacuating-*} \rangle$ states can transition into $\langle \text{assist-secondary} \rangle$; all the others are terminal.

For example, for a commuter to go into $\langle \text{assist-primary} \rangle$ state, they must have high *danger-perception*, high *sang-froid* and high *training*. A complete description of the commuter characteristics, interactions and rules is too long to include in this paper but is available from the authors upon request.

3.4 Variations

For the simulation, there are two places where random variation can be introduced. The first is in the data set generation phase where the concrete data sets are produced from templates. In this scenario, each instance (data set) has a unique, but statistically similar population. The second is the seed used to initialise the random number sequence in the simulation. The major effect of this is to change the exact movement behaviour of hysterical commuters as, for this scenario, the motivation module has been deliberately configured to operate in an essentially deterministic manner. This results from a conscious design decision for the scenario and is not a limitation of the software.

All instances were generated from a single, common template, with changes introduced via a parameter file. The two major parameters altered between variations were *assailant-threat* and *awareness*. The baseline variation was **mid-threat** where *assailant-threat* was 0.75 and *awareness* was randomly selected for each commuter from the uniform range [0.6, 1]. **low-threat** and **high-threat** variations were produced with an *assailant-threat* of 0.5 and 1.0 respectively. Finally, a **high-awareness** variation was derived from **mid-threat**, with *awareness* sampled from the range [0.8, 1]. A single instance of each was generated, from which six runs were performed, varying the random seed for each run.

Four additional instances of the **mid-threat** variant template were generated, giving four distinct, statistically similar populations, each being run once with the same random seed. These are referred to as **random-population-1** through **random-population-4**.

Two structurally different variants of **mid-threat** were also produced. The first was **guided-evac** which made several changes: commuters who were in the $\langle \text{evacuating-at-speed} \rangle$ state and passing through any of nine specified cells were redirected, with 80% probability, to particular exits. In addition, the platform barriers were removed to permit free flow. The intent of this was to introduce a crowd management strategy designed to reduce congestion observed at the exits nearest the platforms. Redirection was performed by adding new induced peep actions to a small number of cells around the platforms.

The second variant, an extension of **guided-evac**, was called **guided-evac-stop-entry**. Here, the arrival of all commuters was compressed into the first 1.2 hours of the simulation, modelling the closure of the station after the blast. Since the number of commuters was not reduced, this also effectively increases the traffic flow in the simulation to an approximation of real-world peak traffic.

4 RESULTS

A summary of results can be found in Table 1.

As expected, commuters in the **high-awareness** runs moved from information-seeking to other post-blast states more quickly. Similarly, the number of commuters in each of the states over time was consistent with the distribution of motivation parameter values.

The simulation results were largely deterministic across runs within a variation. The major random variation was in the exact movement patterns of the hysterical commuters. However, since their presence increases the danger perception of everyone in the same cell, this produced consequent, minor variations in the number of commuters in the various states at any given moment. The most obvious indicator of this was in the number of $\langle \text{assist-primary} \rangle$ commuters, due to their low number. A higher presence of hysterical commuters in a cell pushed trained commuters into $\langle \text{assist-primary} \rangle$, who then gathered more or less secondary helpers based on how far from the blast site they became active. In one of the **random-population** instances, no commuters went into $\langle \text{assist-primary} \rangle$ until close to the end of the simulation. This was due to the random distribution of arrival times of commuters with medical training.

Without more sophisticated evasion behaviour and blast modelling, the number of commuters affected by

Table 1: Selected Results

Variation	Dead/Injured	Evaders	Panicking	Assisting	Evacuating
low-threat	17.0 (0.0)	0.0 (0.0)	507.0 (3.1)	64.0 (8.2)	2187.0 (0.0)
mid-threat	9.0 (0.0)	3.0 (0.0)	502.1 (0.8)	99.3 (3.0)	2187.0 (0.0)
high-threat	17.0 (0.0)	32.0 (0.0)	511.0 (0.0)	52.6 (7.0)	2187.0 (0.0)
high-awareness	14.0 (0.0)	10.0 (0.0)	483.0 (0.0)	90.1 (0.8)	2187.0 (0.0)
random-population	12.75 (2.5)	3.5 (0.6)	496.0 (29.8)	67.0 (44.2)	2187.0 (0.0)
guided-evac	22.0	4	479.0	84.0	2187.0
guided-evac-stop-entry	21.0	11	509.0	207.0	2187.0

Table 1 shows mean and standard deviation (where multiple runs were performed) of selected statistics from the variations, taken at the end of the simulation; the exception is the “Evaders” column which records the number of people the assailant caused to evade over the duration of the run.

the blast cannot be particularly informative. Despite this, the number of people injured or killed seems reasonably consistent with expectations based on the size of the blast. Varying the level of threat the assailant exhibited, with the **low-** and **high-threat** runs, led to a corresponding drop or increase in the number of evading commuters. In the **low-threat** run, no evasion took place.

In all variations, the number of commuters in the station over time increased steadily until the blast occurred, at which point evacuation began. The population dropped rapidly as people evacuated, until most of the remaining commuters were *<catatonic>*, *<hysterical>* or *<assist-*>*, the remainder being those who entered the station subsequent to the IED detonation.

In the **guided-evac-*** variations, the rate of egress from the station was much higher, owing to reduced congestion. It is worth noting that the crowd management strategy of redirecting commuters succeeded with only nine redirection locations used. This suggests that timely placement of a small number of authority figures at appropriate locations can help speed evacuations.

5 CONCLUSIONS AND RECOMMENDATIONS

Despite being proof-of-concept, the results serve to demonstrate a realistic capacity for modelling events of this nature. That a microsimulation-based approach is appropriate for this kind of problem is not new; it is exactly what the approach was designed for. That it can be achieved with meaningful numbers of simulated individuals without recourse to supercomputers or prohibitive runtimes⁷ is, we feel, of some interest. This compares favourably with the agent-based approach of Tsai *et al.* (2011), where the simulation only involved approximately 200 representative individuals in LAX under similar conditions.

The software was not built specifically for this problem. Instead, we have expanded the capabilities of our existing software to accommodate this scenario, without loss of any existing functionality, demonstrating the utility of a compartmentalised approach to modelling. This means that we could equally easily model a terrorist attack involving biological agents rather than explosives.

The work has shown that the motivation system successfully creates different responses to the same stimuli based on internal characteristics. Further modelling will be required to more thoroughly examine the potential of the motivation system, however.

Of particular future interest is the question of whether the evasion caused by threat perception is a factor in reducing the death toll in actual incidents and the extent to which it explains the difference in numbers of deaths between confined and unconfined situations. An answer to this question must be left until trains are added to the model and a more sophisticated assailant design is implemented, including multiple assailant archetypes.

In the near future, we anticipate conducting validation studies using the developed capabilities to model commuter behaviours at the Southern Cross transport hub in Melbourne. Preliminary discussions are also underway with the Los Angeles transport authorities relating to the modelling of fare evasion.

⁷Each run of the scenario took approximately 50 minutes using a single thread of a stock Intel i7 860 CPU.

The work to date has provided us with a simulation environment into which we can now introduce more experimentally derived parameter values, covering not only internal states but also allowing the introduction of passenger schedules and the incorporation of incoming and outgoing trains. The incorporation of additional features, such as: more realistic evasion, security force interdiction, multiple assailants including decoys, caring for the injured, families, biological imperatives, such as hunger, and disorientation during evacuations, is planned for the future.

ACKNOWLEDGMENT

The authors wish to acknowledge the financial support of the ARC via linkage grant LP0882474.

REFERENCES

- Addley, E. (2011, January 12,). 7/7 bus bomber jostled passengers with deadly backpack, inquest told. <http://www.guardian.co.uk/uk/2011/jan/12/77-july-7-bomber-inquest>.
- Cohen, M. L. (1991). Statistical matching and microsimulation models. In *Improving Information for Social Policy Decisions: The Uses of Microsimulation Modeling*, Volume II, pp. 62–85. National Academy Press.
- Connor, R. J., R. Boer, P. C. Prorok, and D. L. Weed (2000). Investigation of design and bias issues in case-control studies of cancer screening using microsimulation. *American Journal of Epidemiology* 151, 991–998.
- Davies, N. (2010). Study of population behaviours in emergency situations. Unpublished Master's Thesis, available upon request.
- Drury, J., C. Cocking, and S. Reicher (2009, March). The nature of collective resilience: Survivor reactions to the 2005 london bombings. *International Journal of Mass Emergencies and Disasters* 27(1), 66–95.
- Georgeff, M., B. Pell, M. Pollack, M. Tambe, and M. Wooldridge (1999). The belief-desire-intention model of agency. *Lecture Notes in Computer Science* 1555.
- Green, A. R., I. C. Piper, D. Keep, R. Bunder, N. Davies, and C. Flaherty (2010). The application of microsimulation to threat modelling. In *National Security Science and Innovation Conference, 9th Safeguarding Australia*.
- Keep, D., R. Bunder, I. C. Piper, and A. R. Green (2011, August). Application of microsimulation towards modelling of behaviours in complex environments. In *Workshop on Applied Adversarial Reasoning and Risk Modeling, AAAI-11*.
- Merz, J. (1991). Microsimulation - a survey of principles, developments and applications. *International Journal of Forecasting* 7(1), 77–104.
- Piper, I. C., D. Keep, A. R. Green, and I. Zhang (2009). Application of microsimulation to the modelling of terrorist attacks. In *Quantitative Risk Analysis for Security Applications Workshop, IJCAI 09*.
- Piper, I. C., D. Keep, A. R. Green, and I. Zhang (2010). Application of microsimulation to the modelling of epidemics and terrorist attacks. *Informatica* 34, 141–150.
- Tsai, J., N. Fridman, E. Bowring, M. Brown, B. Epstein, G. Kaminka, S. Marsella, A. Ogden, I. Rika, A. Sheel, M. E. Taylor, X. Wang, A. Zilka, and M. Tambe (2011, May). Escapes - evacuation simulation with children, authorities, parents, emotions and social comparison. *Proc. 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* 2(6), 457–464.