

**Crowds are made of people:
Human factors in microscopic crowd models**

by

Colin Marc Henein, B.C.S.

A thesis submitted to
the Faculty of Graduate Studies and Research
in partial fulfilment of
the requirements for the degree of

Doctor of Philosophy

Institute of Cognitive Science
Carleton University
Ottawa, Ontario, Canada
August 2008

©2008 Colin Marc Henein

The undersigned hereby recommend to
the Faculty of Graduate Studies and Research
acceptance of the thesis

**“Crowds are made of people:
Human factors in microscopic crowd models”**

Submitted by **Colin Marc Henein, B.C.S.**
in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

Tony White

Tony White, Supervisor

Andreas Schadschneider

Andreas Schadschneider, External Examiner

John Logan

John Logan, Department Chair

Carleton University

August 2008

Abstract

Crowds are physical aggregations of people. There has been an interest in modelling crowds in order to determine the microscopic basis for emergent crowd behaviours. To date, most microscopic models focus more on spatial/physical factors than on human behavioural factors (or vice versa). The isolation of these approaches to crowd dynamics more generally was criticised by Sime, who advocated a more integrated perspective.

We identify microscopic human behaviour simulation, implemented at the heart of microscopic movement models, as a way to act on Sime's concerns. We offer the Microscopic Human Factors methodology to guide this integration: describing (or describing a hypothesis for) a behaviour in the world, taking an individual perspective in describing the essential aspects of the behaviour, and tightly integrating these aspects within an appropriate microscopic formalism.

We present three demonstrations of the methodology. We add to a prominent microscopic model simulations of voluntary pushing (and an accompanying simulation of force and safety), discovery/communication of spatial information, and front-to-back communication (which has not been previously formalised). Key findings of the demonstrations were that (i) limited numbers of agents (modelled individuals) at the rear of a crowd can cause injuries and jamming near the front by triggering and targeting the pushing and leaning forces of others, (ii) knowledge differences lead to behavioural heterogeneity that in highly driven crowds can provoke injuries (that can be resolved through communication) as agents work at cross-purposes; (iii) our conception of front-to-back communication allows agents to dynamically moderate the force in the model to reduce injuries (often at the cost of slower movement). The results of the second demonstration are related to a crowd disaster at Lenin Stadium, while the third demonstration includes a reconstruction of the Who concert disaster at Riverfront Coliseum.

These demonstrations show the practicality and value of the method. Each generates interesting results, discussion, conclusions and future work. Each improves the model relative to Sime's concerns. Each broadens the range of the model, allowing more complex scenarios to be represented. This shows the positive consequences of addressing Sime's concerns in microscopic models: crowds truly are made of people.

*This thesis is dedicated to the memory of
my great-uncle and grandparents:*

Leonard Lisgar Brown (1912–1995)
Andrée Olga Henein (1916–2001)
Edith Katherine Wight (1907–1995)
Murray Ernest Wight (1911–2002)

Acknowledgments

This thesis is the culmination of nine years of work on the Ph.D. — almost five of which have been spent directly on this project. In this time I have had the distinct good fortune to have been supported by many people, both inside and outside the university.

I thank my co-workers at computerActive and Parliant Corporation for their infinite patience with my variable schedule — not to mention their tolerance of my split attention — during work on this project and on the coursework that preceded it. Their support, along with the support of my friends and fellow students (both in Cognitive Science and across the university) has been instrumental in completing this degree. As I cannot name you all I shall name none, but I thank you from the bottom of my heart!

I acknowledge the contribution of the Carleton University Library; the collections, electronic resource and inter-library loans staff have enabled this research. I thank Dr. Keith Still for sending me many crowd videos, and for his informative Crowd Dynamics website. I thank Dr. John Chinneck for his “how to organize your thesis” which was invaluable in structuring this document.

I am grateful to Dr. George Hadjisophocleous, Dr. Guylène Proulx, and the late Dr. Charles Gordon for taking the time to attend the 2004 seminar of an obscure graduate student in a different department, and for taking enough of an interest in my work to guide it, there and then, away from panic and toward human behaviour in crowds. I particularly thank Dr. Proulx for opening her door to me several times through the project, introducing me to important literature including the work of Dr. Jonathan Sime, and for allowing me to audit (read: derail) the lectures in her very valuable course on human behaviour in fire.

I thank Dr. Tobias Kretz, Dr. Andreas Schadschneider and especially Dr. Katsuhiro Nishinari for warmly welcoming an isolated (and timid) modeller into the community of floor field modellers. Not only have I benefited from their interpretations of the work in the field, their encouragement and support at scientific meetings has truly sustained this work — and set a fine example of how science should be prosecuted.

I thank very much the members of the panel that reviewed the proposal for this work: The careful attention of Drs. Andrew Brook, Babak Esfandiari and Robert West improved this project considerably. All urged me to consider the validity of microscopic models, and I particularly thank Dr. Brook in this regard for planting the seed that grew into the methodological part of this thesis (as well as for all he has done, as Chair of Cognitive Science, to facilitate my progress these past nine years). Dr. West suggested several interesting research directions regarding modelling of human behaviour in crowds and — while I regret that more of them did not find their way into this thesis — I hope there will be an opportunity to pursue them in the future.

I especially thank my supervisor, Dr. Tony White, who had the strength to be passionate about this project for five years, even during the inevitable setbacks and worries along the way. His close supervision of the project, careful reading of the results and his research suggestions have greatly improved the work; I have very much enjoyed our stimulating weekly meetings. I am also appreciative of his generosity in funding three trips to present research at distant scientific meetings; he says the money never belonged to him, but I nevertheless appreciate his choice to spend it on this work.

My family has been an important source of support throughout this degree. I especially thank my parents Kringen and Jean-Claude Henein whose unfailing support has been instrumental in reaching this point.

Finally, I acknowledge the immeasurable contribution of my partner, Jen Schellinck, who has sustained and supported me substantially throughout most of this degree. Not only has she directly contributed by participating in our weekly meetings herself and by advising me in matters of research, her sure support in times of stress and her forbearance as I led a double life at work and school has been as invaluable as it has been appreciated. I cannot imagine how I could have finished this thesis and degree without her. Thank you.

Copyright acknowledgments

- Figure 29 is reproduced from (Johnson, 1987b, p. 366, fig. 1, ©1987 University of California Press). Used by kind permission of N. R. Johnson and University of California Press (the latter's permission conveyed through Copyright Clearance Center, Inc.).
- We acknowledge the permission, through Elsevier's general terms of publication, to prepare works derivative of (Henein & White, 2007, ©2006 Elsevier B.V.), excerpts of which appear in chapter 6 of this thesis.
- Portions of chapter 7 are excerpted from (Henein & White, 2006, ©2006 Springer-Verlag Berlin Heidelberg). Reprinted by kind permission of Springer Science+Business Media.
- Portions of chapter 8 are excerpted from (Henein & White, in press-a). Portions of chapter 4 are excerpted from (Henein & White, in press-b). Included by kind permission of Springer Science+Business Media.

Table of Contents

| | |
|---|-------------|
| Abstract | iii |
| Acknowledgments | vi |
| Table of Contents | ix |
| List of Tables | xiv |
| List of Figures | xv |
| List of Algorithms | xvii |
| 1 Introduction | 1 |
| 1.1 The crowd | 1 |
| 1.1.1 Studying crowds | 1 |
| 1.1.2 Crowds are made of people | 3 |
| 1.1.3 People and microscopic models | 5 |
| 1.2 Research objective | 7 |
| 1.3 General approach | 8 |
| 1.3.1 Microscopic Human Factors | 8 |
| 1.3.2 Approach to demonstration | 9 |
| 1.3.3 Demonstrations of value | 9 |
| 1.4 Summary of contributions | 11 |
| 1.5 Overview of thesis structure | 11 |

| | | |
|----------|---|-----------|
| 2 | A brief review of crowds and crowd safety | 14 |
| 2.1 | Force and Injuries in Crowds | 14 |
| 2.1.1 | Concept of levels of service, densities | 14 |
| 2.1.2 | Effects of force: Jamming, Fruin, and Helbing's paradoxes | 16 |
| 2.2 | Selected crowd disasters | 18 |
| 2.2.1 | Bethnal Green Underground Station | 19 |
| 2.2.2 | Lenin Stadium, Moscow | 20 |
| 2.2.3 | Who concert disaster | 21 |
| 2.3 | The myth of panic | 24 |
| 2.4 | Information in Crowds | 27 |
| 2.5 | Summary | 28 |
| 3 | Previous models, critique and problem | 29 |
| 3.1 | Review of previous models | 30 |
| 3.1.1 | Microscopic crowd models | 30 |
| 3.1.2 | Modelling frameworks | 33 |
| 3.1.3 | Models distinguished by their psychological focus | 34 |
| 3.2 | Critique of previous models | 37 |
| 3.2.1 | Microscopic models | 37 |
| 3.2.2 | Modelling frameworks | 40 |
| 3.2.3 | Models distinguished by their psychological focus | 42 |
| 3.3 | Discussion: The problem with microscopic models | 45 |
| 3.4 | Summary | 46 |
| 4 | Microscopic Human Factors | 47 |
| 4.1 | Research objective | 47 |
| 4.2 | Relevant human behaviour | 49 |
| 4.2.1 | The relevance of information processing | 49 |
| 4.2.2 | Behavioural repertoire | 50 |
| 4.3 | The Microscopic Human Factors methodology | 52 |

| | | |
|----------|---|-----------|
| 4.3.1 | The three levels of the methodology | 53 |
| 4.3.2 | Validity and utility of the methodology | 58 |
| 4.4 | Summary | 64 |
| 5 | The floor field model | 67 |
| 5.1 | Cellular automata and Multi-agent systems | 67 |
| 5.2 | Design of model | 69 |
| 5.2.1 | Physical environment | 69 |
| 5.2.2 | Floor fields | 70 |
| 5.2.3 | Cell selection | 74 |
| 5.3 | Sensitivity parameter effects | 74 |
| 5.4 | Model evolution | 75 |
| 5.5 | Summary | 76 |
| 6 | Voluntary pushing, force and crowd safety | 78 |
| 6.1 | Specification: Forces in crowds | 80 |
| 6.2 | Reduction: Local rules for force behaviours | 81 |
| 6.3 | Implementation | 82 |
| 6.3.1 | Force Field Addition | 83 |
| 6.3.2 | Changes to agent behaviour | 84 |
| 6.3.3 | Time step Progression | 86 |
| 6.4 | Results: Adapting the model for force | 88 |
| 6.4.1 | Model Parameters | 89 |
| 6.4.2 | Neighbourhood size | 89 |
| 6.5 | Results: Effects of force | 91 |
| 6.5.1 | The effect of k_S on agent drive | 93 |
| 6.5.2 | Agent density and the aisle | 95 |
| 6.5.3 | The aisle and model realism | 99 |
| 6.6 | Results: Comparison of measures | 100 |
| 6.7 | Results: Effects of injuries | 102 |

| | | |
|----------|---|------------|
| 6.7.1 | Mechanism underlying injury effects | 105 |
| 6.8 | General discussion | 107 |
| 6.9 | Model validity | 111 |
| 6.10 | Summary | 113 |
| 7 | Discovery and communication of spatial information | 116 |
| 7.1 | Specification | 117 |
| 7.2 | Reduction | 118 |
| 7.3 | Implementation | 120 |
| 7.3.1 | Representations of space | 121 |
| 7.3.2 | Physical discovery of spatial information | 122 |
| 7.3.3 | Communication of spatial knowledge | 124 |
| 7.3.4 | Algorithm | 125 |
| 7.4 | Results | 125 |
| 7.4.1 | Homogeneous crowd | 128 |
| 7.4.2 | Heterogeneous crowd | 131 |
| 7.4.3 | Communication in a heterogeneous crowd | 137 |
| 7.5 | Discussion | 140 |
| 7.6 | Model validity | 142 |
| 7.7 | Summary | 144 |
| 8 | Front-to-back communication | 147 |
| 8.1 | Specification | 148 |
| 8.2 | Reduction | 151 |
| 8.3 | Implementation | 153 |
| 8.3.1 | Algorithm for front-to-back communication | 155 |
| 8.4 | Laboratory scenario | 155 |
| 8.4.1 | Effect on injuries and exits | 157 |
| 8.4.2 | Effect on the aisle | 159 |
| 8.5 | Who concert disaster scenario | 160 |

| | | |
|-----------|---|------------|
| 8.5.1 | Reconstructing the plaza at Riverfront Coliseum | 161 |
| 8.5.2 | Analysis of the disaster | 163 |
| 8.5.3 | Comparison with front-to-back communication | 165 |
| 8.6 | Return to Lenin Stadium | 168 |
| 8.6.1 | Results | 169 |
| 8.7 | Model validity | 171 |
| 8.8 | Discussion | 173 |
| 8.9 | Summary | 175 |
| 9 | General discussion | 177 |
| 9.1 | Crowd disasters | 177 |
| 9.1.1 | Training in front-to-back communication | 177 |
| 9.1.2 | Often there is plenty of time to support the outcomes we want | 179 |
| 9.1.3 | Was it a failure of front-to-back communication? | 181 |
| 9.2 | The Microscopic Human Factors methodology | 183 |
| 9.2.1 | The value of the specification | 183 |
| 9.2.2 | The value of the reduction | 184 |
| 9.2.3 | Effect on implementation | 185 |
| 9.3 | Sime's concerns | 186 |
| 10 | Conclusions and future work | 189 |
| 10.1 | Conclusions | 189 |
| 10.2 | Contributions | 191 |
| 10.3 | Future work | 192 |
| 10.4 | Closing remarks | 194 |
| | References | 196 |

List of Tables

| | | |
|---|---|-----|
| 1 | Summary of MHF methodology levels, their content, purpose and value in validation | 64 |
| 2 | Parameter table with default settings | 89 |
| 3 | Cell transition probability with no occupied cells | 94 |
| 4 | Cell transition probability with forward neighbour occupied | 98 |
| 5 | Comparison of FFM and force analyses | 109 |

List of Figures

| | | |
|----|--|-----|
| 1 | Arched jamming at exits | 16 |
| 2 | The three layers of the MHF methodology | 54 |
| 3 | A trivial example of the Microscopic Human Factors methodology at work. | 55 |
| 4 | Comparison of state and movement in a cellular automaton and a multi-agent system. | 68 |
| 5 | Floor field model space with static field | 71 |
| 6 | Floor field model space with dynamic field | 73 |
| 7 | Non-force changes to the floor field model | 90 |
| 8 | Time steps to complete exit with and without force | 92 |
| 9 | Crowd density schematic | 93 |
| 10 | Aisle evolution | 96 |
| 11 | Number exiting in fixed time with and without force | 101 |
| 12 | Number injured and exited in fixed time with force and injuries | 104 |
| 13 | Standard deviation of number exiting in figure 12 | 105 |
| 14 | Large force block with many injuries | 107 |
| 15 | Devolution of the aisle | 108 |
| 16 | Waves of force | 112 |
| 17 | Discovery field | 123 |
| 18 | Physical environment for discovery and communication model scenarios | 127 |
| 19 | Agents exiting and injured with ample exits | 129 |
| 20 | Agent egress measures in a homogeneous crowd with plenty of exits ($\phi = 125$) | 130 |
| 21 | Static and discovery fields in heterogeneous scenario | 131 |

| | | |
|----|--|-----|
| 22 | Three groups of agents at discovery of blocked exits | 133 |
| 23 | Agents injured in exiting from the blocked exit scenario without communication | 134 |
| 24 | Exiting agents in the blocked exit scenario without communication | 136 |
| 25 | Injuries with communication enabled | 138 |
| 26 | Exiting agents in the blocked exit scenario with communication | 139 |
| 27 | Effect of front-to-back communication: laboratory scenario | 158 |
| 28 | Disappearance of the aisle region with front-to-back communication | 159 |
| 29 | Plaza at Riverfront Coliseum | 162 |
| 30 | Riverfront Coliseum scenario without front-to-back communication | 164 |
| 31 | Perpendicular forces create injury hot spot at Riverfront Coliseum | 166 |
| 32 | Effect of front-to-back communication at Riverfront Coliseum | 167 |
| 33 | Front-to-back and spatial communication in a barrier formation | 169 |
| 34 | Comparison of the effect of communication types on a barrier formation | 170 |

List of Algorithms

| | | |
|---|--|-----|
| 1 | Floor field model time step function | 77 |
| 2 | Step function augmented by pushing and force additions | 87 |
| 3 | Step function modified for discovery and communication | 126 |
| 4 | Step function modified for front-to-back communication | 156 |

Chapter 1

Introduction

1.1 The crowd

Crowd situations are one of the few situations in which humans flock together. Although most crowds we encounter in our daily lives are safe, media reports remind us that different types of crowd situations can be dangerous. Evacuations in fires, for example, can become competitive and people can be badly injured or killed through crowd forces just as they can be overcome by smoke and flames (Chertkoff & Kushigian, 2001). Acquisitive crowd disasters, sometimes called *crazes*, can occur when people compete to obtain a limited resource (like food, water, or one of a limited number of gatecrasher specials at a retail outlet). Large numbers of people moving in a constrained environment, even if they have the same goals and are not competitive, can also find themselves in dangerous conditions; at the Hajj, for example, crowd disasters have occurred as 2.5 million pilgrims attempt to complete religious obligations, but also as pilgrims simply return to their accommodations after the completion of the ritual (Crowd Dynamics, 2007).

1.1.1 Studying crowds

Crowds are aggregations of humans. These aggregations themselves seem to have certain behaviours and properties (e.g. jamming at bottlenecks) that *emerge* from the collective behaviour of individuals within the crowds. Like other aggregations of animals (e.g. cattle herds, fish schools) and material (e.g. fluid and particle flow) that have their own emergent dynamics, there has been an interest in determining the basis of crowd-level effects and properties, such

as in cases where large groups of people must move through a doorway, narrow hall or other restricted environment. If we are armed with an understanding of how collections of individuals create emergent crowd-level behaviours, we may be able to create spaces or suggest behaviours that promote safe and efficient movement of crowds.

As crowd behaviours arise from large groups of people, and as crowd behaviours of interest may be dangerous, computer modelling has emerged as a key tool to study the foundations of crowd behaviour and to predict crowd movement through physical environments. These models explore scenarios as disparate as planning designated walking routes for large urban festivals (Batty, Desyllas, & Duxbury, 2003), interactions between military and civilian groups in wartime (Nguyen, McKenzie, & Petty, 2005), estimating time for people to exit from a high-rise building (Fahy, 1994), and many others. The purpose of the models can be practical (as are the models just described) or more theoretical — considering factors that underlie particular patterns of movement or crowd effects. Additionally, different approaches to modelling exist, with some models being based on traditional numerical-continuous techniques adapted from physics (e.g. Helbing, Farkas, Molnár, & Vicsek, 2002) and others being more based on discrete modelling techniques like cellular automata (e.g. Burstedde, Klauck, Schadschneider, & Zittartz, 2001).

Our interest lies in the area of microscopic modelling of crowds. Microscopic models are designed to work just like real crowds. They directly simulate the behaviour of each individual in a crowd, with collective effects emerging from individual behaviours. They ask questions like *which* individual actions lead to the complex collective behaviours, and *how* these individual actions combine to create crowd behaviours. The aim of this type of research is often qualitative, studying the behaviour of abstract crowds, rather than producing simulations that aim to quantify and predict crowd behaviour in specific environments. In short, microscopic modelling looks at the mechanism for emergent crowd effects.

The general microscopic strategy is to model individuals using simple rules for behaviour. These rules should be *local*: Like real people in real crowds, agents¹ perceive only their immediate surroundings. These simple local rules determine the agent's behaviour, and — when collections of agents are gathered together — emergent behaviours occur and can be studied at the level of the crowd. The emergent behaviours come directly out of the simple rules. This offers the opportunity for experimentation upon the mechanisms of the model. Through this experimentation causal connections can often be established between the rules (at the level of the individual) and the effects (at the level of the crowd). As most microscopic crowd modelling has taken place within the domain of statistical mechanics in physics, primary interest has been in modelling physical and movement effects.

1.1.2 Crowds are made of people

Unlike cattle, fish and seeds in a hopper, what unifies all crowds — disastrous or no — is the fact that they contain a number of humans situated in an environment. Accordingly, in considering crowd behaviour we must remember that we are studying a phenomenon of groups of social actors driven by human behaviour. Disregarding this while trying to understand crowd disasters can lead to over-emphasis on such factors as the physical environment, and over-generalisation of crowd disasters to physical domains such as particle and fluid dynamics. In this regard, it is worth noting that although many non-cognitive factors are legitimately implicated in crowd disasters (e.g. those described in the opening paragraph of section 1.1.1), such disasters never arise in the analogous inanimate situations. There may be some physical similarity between the arrangements of seeds in a hopper, particles in a hydrodynamic system, and persons in a crowd, for example, but people have many capacities that the other systems do not. For example, people have the capacity to extract information from the environment, alter their goals and change direction, communicate verbally and non-verbally and apply force at their discretion. Although

¹In this thesis we shall use the term *person* to refer to humans in real crowds. We shall use the term *agent* to refer to a modelled individual in a simulated crowd. Crowd, crowd member and individual are used interchangeably to refer to agents or people, the meaning should be clear from the context; where necessary we specify that we are talking about a *real* crowd as opposed to a modelled one.

we must always remember that a crowd is a collection of people, it is commonplace for us not to do so.

Although this point may seem evident, Jonathan Sime, an environmental psychologist, cautioned building designers against creating what he termed *ball-bearing designs* that treat people “as if they are non-thinking objects rather like the elements of a building structure” (1985). To psychologists and engineers, he pointed out that this kind of thinking is encouraged by arbitrary distinctions between academic disciplines; while engineers create ball-bearing designs, psychologists divorce crowds from their environment:

Psychology and engineering are characteristically mutually exclusive in their focus on the perspective of crowd members who *think* and *behave* (psychology) or on static and dynamic objects (engineering) . . . The history of psychology, and indeed crowd psychology, is one in which the physical environment in which behaviour takes place has been ignored. While this might justifiably frustrate engineers, it is equally the case that social scientists are bemused by the lack of attention in engineering to the psychological and social aspects of crowd safety. . . (1995)

Let us take an example to illustrate the necessity of studying human behaviour in crowds. Many crowd disasters have occurred when people at the rear of a crowd have an expectation of what is going on at the front that leads them to engage in counter-productive behaviours (like pushing). This situation, for example, may arise after a fall on stairs when the front of a crowd stops to allow the fallen person to rise. Those at the rear of the crowd — unaware of the reason for the delay — press forward. The combined forces of the crowd cause a pile-up and in extreme circumstances can lead to death by crushing for those at the focus of the force. This situation can be summed up by stating that “crowd incidents often exhibit . . . a failure of front-to-back communication” (Pauls, 1984; see also Fruin, 1993).

In trying to prevent crowd disasters we might accordingly try to increase the safety of stairs to prevent falls (an engineering approach). It is worth stopping to consider, however, whether the situation might equally have been averted by alerting those at the rear of the delay, its cause and its transience (a behavioural approach). While we might say that those at the rear caused the disaster through their pushing behaviours (a physical characterisation), we can equally say that the disaster had its origins in human factors: ignorance (of the fall), assumption (of the

conditions at the front), expectation (of the normal progress of the crowd), etc. As Sime suggests, we might expect that researchers studying pedestrian dynamics would take both physical and human factors into account, but frequently this does not occur.

Another example of a lack of attention to human behaviour is the persistence of *panic* as a description of behaviour in crowd disasters. Despite the fact that general panic as a conception of crowd disaster behaviour has been academically discredited (see section 2.3 for more on this point), its constant invocation by news reporters (to say nothing of fantastic images provided by the entertainment industry) continues to fuel belief in this construct. Besides being a non-explanatory construct, the invocation of panic to explain crowd disasters tends to shift attention away from designers, planners and responsible authorities; safety systems designed by these experts, the argument goes, can hardly be blamed if people ‘lose their heads’ and behave in irrational ways (Sime, 1995). Panic, although discredited, is a compelling concept that persists in crowd research.

1.1.3 People and microscopic models

But what of Sime’s concerns regarding the artificial separation of engineering and psychological research when it comes to the *modelling* domain? Sime did not take a position on microscopic models in particular, although he did comment in 1995 on models that study fires and similar dangers in buildings and ships. He reported that modellers recognised the problem of omitting human cognition, decision making and social behaviour from models, but were not addressing it due to the difficulty of implementing cognitive simulations, saying: “it is easier to model people as homogeneous objects than interacting individual and group decision makers.” As evacuation studies have not found identical automatons to be a good model of people, however, Sime suggests greater psychological relevance is required for accurate modelling in these circumstances.

Microscopic models have evolved significantly in the intervening years. Given the focus of these models on the behaviour of individuals, and the advances in both modelling technique and computational power, we might expect that human behaviour would now play a significant role

in modelling the behaviour of individuals. Unfortunately, this is not the case: the issues raised by Sime persist. Microscopic models are generally identifiable as taking either an engineering or psychological perspective rather than the necessary integrated view, and we are unaware of a model that integrates these two aspects of crowd behaviour satisfactorily (see section 3.2 for more on this point).

Confirming Sime's ball-bearing concerns, engineering-heavy models suggest crowds are simple and behaviourally homogeneous — that is they are composed of interchangeable automata that react identically to a given stimulus and do not maintain memory or experience that makes them different from a neighbouring individual. This strategy produces simplistic results that are unrealistic when simulation results are compared to real crowd scenarios; these often include, for example, varying information, changing beliefs, and communication. Modelling homogeneous agents denies the contribution of information, varying goals, and many other human characteristics to crowd behaviour.

The models that do take account of human behaviour tend to suffer from missing or poor integration with physical factors. Either they isolate themselves by focusing on behaviour to the detriment of a physical simulation, or they provide a weak integration of behavioural and physical simulation leading to uncertainty as to the mechanism of emergent crowd effects. As real-life crowd situations involve both human and physical factors, it is our contention that the lack of appropriate models is an impediment to understanding crowd dynamics.

An additional and related problem is that microscopic modellers on both sides of the disciplinary divide continue to model panic behaviours. This may be due to a lack of familiarity with arguments against this view. Alternatively it may be because an assumption of panic significantly reduces the work of creating a model.

These two problems — a continuing lack of interdisciplinary integration in microscopic modelling, and a continuing focus on panic as a driver of behaviour — are the rationale for our research.

1.2 Research objective

The central theme of this thesis is that the interaction between behavioural and physical factors in crowd dynamics requires that these be treated in an integrated fashion to obtain a model that can reasonably be applied to people.

As current microscopic models fail to satisfactorily integrate these factors, we consider how they *can* be integrated in the context of a microscopic model of crowd safety and dynamics. The research problem became evident while taking a cognitive science perspective on the state of microscopic modelling and grew out of our concerns that existing microscopic models displayed unrealistic behaviour. We developed this position independently, and it was strengthened when we uncovered Sime's statements of his similar concerns (1985, 1995). We go beyond Sime, however, noting that he restricted himself to encouraging modellers to improve their own models; this thesis proposes a methodology for incorporating human factors into microscopic crowd models.

Research objective: *To develop a methodology to integrate relevant human behaviours with physical factors at the heart of microscopic crowd models for the purpose of improving relevance to people and studying emergent crowd-level effects.*

In this thesis, we develop a methodology that will allow for Sime's concerns to be addressed by implementing human factors in a microscopic model. A key aspect of this task is to implement these factors without violating the integrity of the microscopic method (simulation of emergent behaviour via collective action of individuals following local rules); microscopic models can support integrated simple rules that together encompass *both* a behaviour simulation and a physical one. By keeping the focus on the abstraction, simplicity and integration of rules, we preserve the ability to experiment on the model to relate emergent crowd-level effects and the mechanism that underlies them (we shall return to this point in chapter 4).

The value of the methodology is then demonstrated through its use in integrating three human behaviours into an existing microscopic model, and analysing the effects of these behaviours on crowd dynamics and safety.

1.3 General approach

1.3.1 Microscopic Human Factors

We propose the *Microscopic Human Factors methodology* (chapter 4) to meet our objective within the context of microscopic modelling. The methodology involves three complete descriptions of a human behaviour of interest at increasing levels of abstraction: the level of real behaviour in real crowds, a description of the essentials of this behaviour (phrased in the language of individuals), and an implementation in an appropriate microscopic model's formalism. This methodology is intended to generate an abstract model of the human behaviour of interest, making explicit the modeller's theoretical stance toward the essential mechanism of the behaviour, and producing a representation of this mechanism that can be tightly integrated with a microscopic model. By requiring that the behaviour be implemented at the heart of the microscopic model's simple rule set, the methodology helps to ensure that experimentation upon the implemented model can reveal the effects of the human behaviour.

Ethical concerns preclude the experimental testing with human participants of many crowd hypotheses, particularly those involving force and crowd safety. The abstract focus of the microscopic model provides an important platform from which to evaluate how various factors can interact in crowds. The scope of this project does not allow for development of non-microscopic models that would provide quantitative predictions in specific scenarios, and we do not aim to provide models that replace direct experimentation. Notwithstanding this, a microscopic model incorporating human factors, such as may be generated using the methodology, provides an environment within which to study the plausibility of crowd safety hypotheses and to examine factors that may influence their effects. As with much research in the social sciences, we stress the importance of confirming evidence from other disciplines and case studies. We shall return to the question of validity in more detail in section 4.3.2.

We have chosen to implement three human behaviours as microscopic human factors to demonstrate the effectiveness of the methodology and the consequent improvement (as per Sime's concerns) of a microscopic model. In the next section we describe the environment

in which these demonstrations are carried out, and then we will present the details of each demonstration.

1.3.2 Approach to demonstration

Several different modelling approaches have been put forward for computer modelling of crowds, including differential equation models, numerical attraction-repulsion models and cellular automata. These types of models all have their strengths and weaknesses, but the principles of multi-agent systems and swarm intelligence are a natural fit for microscopic crowd models. These approaches emphasise the simulation of a complex system by modelling of a population of independent agents (e.g. Bonabeau, Dorigo, & Theraulaz, 1991). Each agent is guided by simple rules that consider information available locally in determining the agent's subsequent behaviour. Agent state is portable. This is analogous to the information resources available to an individual in a large crowd; persons within crowds do not have access to an overall perspective of the crowd's conformation and must act with only their memory, experience and immediate local conditions as sources of information.

We have adopted the floor field model (Burstedde et al., 2001; as modified in Kirchner & Schadschneider, 2002) as the base for our microscopic, multi-agent model. This cellular automaton model was well-suited for conversion to a multi-agent system due to its discrete grid environment and emphasis on local processing by agents.

1.3.3 Demonstrations of value

To demonstrate the Microscopic Human Factors methodology, we have added three behaviours to the floor field model: voluntary pushing, discovery and communication of spatial information, and front-to-back communication. The first demonstration, of voluntary pushing (chapter 6; Henein & White, 2005, 2007), also includes the addition of a simulation of the effects of force on agents; the voluntary pushing is a human behaviour, while the corresponding force simulation allows the model — originally restricted to modelling agent movement times — to consider crowd safety through injuries as a dependent variable.

Retaining the voluntary pushing and force simulation, the demonstration of discovery and communication of spatial information adds the ability for agents to have a time-evolving representation of the environment, and to communicate this outlook to others (chapter 7; Henein & White, 2006). By allowing agents to change their view of the world, and by allowing them to communicate these changes to others, the model is able to represent human capacities that the simpler particle-based floor field model can not capture. The crowd heterogeneity that results from simulating agents with multiple views of the world can also cause groups of ‘like-minded’ agents to form. We study the interactions between these groups in various crowd conditions.

In order to show the methodology can also operate in a more prospective mode, we build on the results of the first two demonstrations to study locally-mediated front-to-back communication in crowds (chapter 8; Henein & White, in press-a). As noted in section 1.1.2, failure of front-to-back communication is commonly described in crowd disasters (Pauls, 1984; Fruin, 1993). The mechanisms for successful front-to-back communication, however, and its ability to moderate crowd forces have not been previously studied. We examine the effects of such communication in three scenarios: a laboratory-type scenario allows comparison with the results of the previous two chapters, while the other two relate to real crowd disasters, including a reconstruction of the Who concert crowd disaster (see section 2.2.3). In the case of the Who disaster, we qualitatively examine the forces in the original disaster, and then evaluate the benefits of front-to-back communication.

These three demonstrations have been selected in order to demonstrate the methodology at work toward three different ends. The first and second applications demonstrate well-documented crowd behaviours, while the third investigates a behaviour that has been frequently cited, but has not been previously studied in depth. The first and third applications demonstrate changes where the goal is to add human behaviours to agents and then investigate the emergent effects that ensue, while the second application demonstrates changes intended to directly create groups of different-behaviour agents and then examine the interactions of these groups.

For each of these demonstrations we will provide a qualitative evaluation of the resulting changes in crowd dynamics and safety, and discuss how the rules introduced by the methodology underlie the emergent crowd effects produced. This will demonstrate the practicality and

utility of the methodology in creating models that generate interesting results, discussion, conclusions and future work. By demonstrating the additional value in representing human factors in microscopic models, we will ultimately support Sime's point regarding the importance of integrating human and physical approaches to crowd dynamics.

1.4 Summary of contributions

This thesis makes two kinds of contribution to knowledge (we provide a list of specific contributions in section 10.2). First, the Microscopic Human Factors methodology is offered as a way to take account of Sime's concerns by bringing human factors into the heart of microscopic models. This methodology provides guidance in improving microscopic models to respond to Sime's criticisms while preserving the unique contributions of the microscopic approach.

Second, our demonstrations of the methodology, and the associated qualitative analyses, each make contributions as they explore the mechanisms of the behaviours on which they focus, and discuss their effects on movement and safety in crowds. This brings increased understanding of the phenomena being demonstrated. Each demonstration is a contribution, as well, to the floor field model's research community, adding three different human factors into the context of this microscopic model. In addition, the first demonstration provides a microscopic simulation of crowd forces and safety, previously missing from this model.

1.5 Overview of thesis structure

In this thesis we divide the material presented into two categories. The first category presents work that has come before: background information, and a review of the literature. The second category presents our work and original research. The dividing line between these two categories occurs midway through chapter 3, at the point where it turns from a description of the literature of microscopic models to a critique of these models. In chapter 4, with a foundation in crowd dynamics and the literature, we define our research objective and begin in earnest our contributions with a discussion of the Microscopic Human Factors methodology.

Beginning with chapter 2, the document proceeds as follows:

- Chapter 2 provides background information on pedestrian behaviour, crowd dynamics and crowd safety; three short case studies are presented, and a discussion on information in crowds. The purpose of this chapter is to give the reader enough of a background on what is known about people in crowds to motivate the decisions taken during the modelling process. Additionally, the case studies (a traditional way to study the human element in crowds) will be referred to frequently in the remainder of the thesis.
- Chapter 3 begins with a presentation of the literature to date on human behaviour in microscopic crowd modelling. (From here, we begin to focus on our own work, ideas and research.) Having described the previous work, we then turn to a critique of the models presented, identifying why — despite their benefits — they fail to satisfactorily integrate human factors and physical factors together in the context of a microscopic model. The problem of our inability to find a satisfactory model, and our discussion of the importance of this integration is the topic of the closing section of the chapter; the goal of this key section (3.3) is to succinctly motivate the research in the rest of the thesis.
- Chapter 4 begins by explicitly stating the research objective that defines the thesis, inspired by the critique and problem summarised in section 3.3. We then turn to our approach to reach this objective. We describe the Microscopic Human Factors methodology, how to build a model using it, our arguments for establishing validity in models created with it, and the limits that can be expected of this validity in an abstract microscopic model.
- Chapter 5 is a brief return to the work of others. As the methodology requires tight integration of human factors with existing microscopic models, we need to describe such an existing model in sufficient detail to make clear the process of integration. This chapter provides the details of the floor field model (c.f. Burstedde et al., 2001; Kirchner & Schadschneider, 2002), which we shall use as the base for the demonstrations to follow.

- Chapters 6–8 demonstrate the Microscopic Human Factors methodology in action. As discussed above, three experiments have been carried out to show how human behaviours can be modelled microscopically, and the benefits that can be obtained from doing so.
 - Chapter 6 focuses on the development of a model of voluntary pushing, and includes a simulation of force and crowd safety.
 - Chapter 7 builds on chapter 6 to add discovery and communication of spatial information.
 - Chapter 8 builds on chapter 7 to study front-to-back communication in a laboratory-type scenario, and in two scenarios related to real-world crowd disasters.
- Chapter 9 is a short general discussion. It begins by bringing together the crowd dynamics results obtained from the three separate demonstrations to draw joint conclusions, relating these to research by others in this field, and suggesting future work inspired by this discussion. It returns to the methodology itself, making observations from the demonstration process about the methodology's benefits. Finally, the chapter returns to our original motivation of addressing Sime's concerns within the microscopic model context, arguing that the methodology has improved microscopic modelling by taking Sime's concerns into account.
- Finally, chapter 10 summarises the conclusions obtained from the work, the contributions made by the thesis, presents future work and some brief concluding remarks.

Chapter 2

A brief review of crowds and crowd safety

In order to build good crowd models it is necessary to understand how real crowds behave. Traditionally this understanding has come through the study of case studies, but much is also known about crowds from a theoretical perspective. By reviewing the human element as well as the mechanical element of crowd behaviour we will be in a position to evaluate existing crowd models and determine what must be added to them to improve our understanding of crowds. We will briefly describe the basics of pedestrian dynamics and force in crowds, then describe three short case studies, the myth of panic and conclude with a discussion of information in crowds.

2.1 Force and Injuries in Crowds

2.1.1 Concept of levels of service, densities

Pedestrian dynamics are in some respects similar to traffic flows. Both pedestrians and vehicles can exhibit following behaviours, can move through space that is set out specifically for their use, can group themselves (or are grouped) into lanes that exhibit smooth flow behaviours. Pedestrians, however, have much more flexible behaviours than vehicles do, and these more flexible behaviours place them into a separate class (Blue & Adler, 2001). Pedestrians:

- can turn and reverse direction almost immediately at normal walking speeds
- enjoy unregulated behaviour compared to the strict controls for vehicle traffic

- are free to vary their speed and can accelerate to full speed from a standstill
- are free to occupy any part of a walkway
- often walk side-by-side or in clusters

Moreover, the acceptable dynamics for pedestrians are different from vehicle dynamics. For example, it has been noted that “while at the immediate and superficial level encounters on the street are hardly noticeable and devoid of pleasantries and warmth, pedestrians do, in fact, communicate and do take into account the qualities and predicaments of others in regulating their behavior” (Wolff, 1973). Thus, as pedestrians are similar but not identical to vehicles, modelling and planning for pedestrians requires pedestrian-specific approaches and techniques.

Planning for pedestrians requires a knowledge of the capacity of an environment to support pedestrian behaviours, and also the number of pedestrians that will occupy a space. Fruin has created the concept of levels of service in pedestrian environments, relating pedestrian density to safety and possible movement rates (in the case of walkways) and psychological stress (in the case of queueing). Fruin stresses that these human considerations are very important to take into account in designing public spaces (1971).

Physical aspects of the environment will affect the density of pedestrians, and can affect crowd distribution dramatically. For example, if people must move through a bottleneck, the density of a crowd that was diffuse can increase dramatically. It has even been shown that when pedestrians move through a bubble of expanded space (such as an open hall) in a corridor of otherwise constant size, they experience an increase in density; even though the number of persons was appropriate for smooth flow in the corridor, the spreading out in the expanded area creates a bottleneck and increased density when the crowd must reform into the corridor width. Density is also increased where two large flows of pedestrians must cross, disrupting the smooth flow (Still, 2000). Thus, event planners, engineers and architects must carefully consider how spaces are designed and used with an eye to ensuring pedestrian safety — even under extreme conditions.

2.1.2 Effects of force: Jamming, Fruin, and Helbing's paradoxes

Studying the dynamics of crowd situations requires as much consideration of crowd forces as it does of crowd movement and behaviour. Fruin states that forces generated within crowds are caused by individuals physically pushing others, or through forces transmitted during leaning on — or falling against — other people. There is a domino effect by which force originating with one person can be transmitted and augmented through the next person, and this additive effect can be the source of waves of force and serious safety issues (Fruin, 1993).

Moderate forces tend to cause inefficiencies in movement at bottlenecks, like doors for example. Individuals tend to form a rainbow-like structure, resembling concentric semi-circular arches centred on the doorway (see figure 1). Studies of this phenomenon indicate that the arches have some structural stability and when loaded from behind tend to withstand a certain amount of force. At a certain point, the innermost arch collapses towards the door (resulting in the exit of a few people immediately in front of the door). The arch immediately behind the now-exited front rank moves forward, and jams due to the fact that everyone in the now front-most arch presses in toward the door. The cycle repeats with a build-up of force behind this jammed arch. Video footage of exits reinforces the fact that this somewhat idealised view can be further hampered by shoving that sets up oscillatory forces lateral to the direction of travel through the door that also disrupt movement through the exit. Crowd researchers tend to label these effects *jamming* or *clogging* at a bottleneck.




Figure 1: Arched jamming at exits. Left configuration is stable until front (white) arch breaks. Some agents from next rank (light grey) move forward (middle configuration) while white arch exits, forming a new arch (right configuration).

When forces build up to the point that they can no longer be described as moderate, serious crowd situations can cause injuries and fatalities. Pedestrians are capable of packing themselves

into a space well beyond what an environment will sustain and, by doing so, an individual gives up control over their movement in a situation. Fruin has concluded that this point occurs around densities of seven persons per square metre:

It is difficult to describe the psychological and physiological pressures within crowds at maximum density. When crowd density equals the plan area of the human body, individual control is lost, as one becomes an involuntary part of the mass. At occupancies of about 7 persons per square meter the crowd becomes almost a fluid mass. Shock waves can be propagated through the mass sufficient to lift people off of their feet and propel them distances of 3 m (10 ft) or more. People may be literally lifted out of their shoes, and have clothing torn off. . . (1993)

It is when these forces pin individuals against immobile objects — for example a wall, obstacle, or even another person being pushed in a different direction — that injury or death occurs. Fruin reports that the most common cause of death in crowds is compressive asphyxia — occurring when the pressure on a person's chest exceeds the capacity of their chest muscles to expand the lungs. By analysing bent steel railings after fatal crowd disasters it is known that crowd forces can exceed 4500 N (Fruin, 1971). Thus, it is vital to maintain crowd safety that pushing be kept to a minimum in situations of high crowd density, and that crowd densities be kept below critical levels.

Helbing has studied patterns of behaviour that occur when crowd forces reach extreme levels. Two of his conclusions that are important to understanding crowd force behaviour are expressed as paradoxical statements: 'faster is slower' and 'freezing by heating' (Helbing et al., 2002).

In 'faster is slower', the observation is made that slow and leisurely patterns of movement promote smooth flows and faster exits. By increasing the speed, and therefore the energy and momentum in the system, the flow becomes uncoordinated and less efficient. Consider a room full of theatre-goers exiting after a long evening. Helbing's thesis predicts that if they move slowly but steadily toward the exit, then although the exit may take some time it will be smooth and unmarked by clogging and jamming. By contrast, consider a crowd of shoppers anxious to obtain one of a limited number of so-called door-crasher specials. If these individuals move quickly from their cars to the entrance, higher speeds may cause a jam at the door. Even without considering the possibility of people actively pushing that causes crushing, the retail situation is

much less safe with an increased risk of jostling, stumbling and falling behaviours due to the higher speeds of individuals. The net result, according to Helbing is slower collective progress due to higher individuals speeds.

In 'freezing by heating' the observation is made that when speeds and energies are low, crowd tolerances (in a physical sense) may be higher than when speeds and energies are high. Consider two facing groups of pedestrians crossing the street. These two groups need to pass through one another in the crosswalk. Slow speeds promote smooth flow as people have time to find a space through the opposing crowd. If the groups attempt to cross at a run, however, there will be insufficient time to make this accommodation and there will be collisions and jamming in the crosswalk. With sufficiently large crowds and high energy and determination this kind of situation can become a crowd disaster as one crowd pushes against the opposite to get through and vice versa (an example of this kind of disaster will be found in section 2.2.2). Thus, Helbing concludes that slower collective progress can result from higher individual energy in crowds.

A study of footage at the Hajj in 2006 documented the kind of shockwaves described by Fruin (Helbing, Johansson, & Al-Abideen, 2007). As densities and forces increased during the stoning of the devil ritual, the crowd moved from a laminar flow to stop-and-go movement to shockwaves with a kind of incoordination reminiscent of air turbulence. Force built up over time, and was suddenly released in a manner reminiscent of earthquakes.

In summary, the study of pedestrian and crowd behaviour is different from vehicle flows and from the study of fluid dynamics and particle dynamics. Pedestrians' differing physical capabilities are one reason, but crowd pressures due to the desire to achieve certain goals (like passing through a doorway) are another significant factor in explaining crowd behaviour. Crowd forces override individual movement decisions when they become intense, can reach prohibitive pressures, and in extreme situations cause death through compressive asphyxia.

2.2 Selected crowd disasters

The purpose of this section is to present three crowd disasters, in order to put the preceding discussion and the introduction into some context. Additionally, these disasters will provide

examples that can be referred to later. Only the principal details of each disaster are presented here, with slightly more detail in the Who concert tragedy.

2.2.1 Bethnal Green Underground Station

On the evening of March 3, 1943, in the London neighbourhood of Bethnal Green, air raid sirens sounded due to the possibility of German reprisals for an earlier raid by allied forces. In fact, no bombs were dropped, but media reports suggest that loud experimental British anti-aircraft missiles were tested (Lack, 2003). An inquiry report at the time suggested that “a salvo of rockets was discharged from a battery a third of a mile away” (Dunne, 1943/1999). A survivor of the incident recently recalled that “everybody surged forward, shouting, “There’s bombs, there’s bombs, get down”” (Lack, 2003). The inquiry report mentions that “there were some cries reported that ‘they were starting dropping [the bombs], that it was a land mine, and other alarming observations”” (Dunne, 1943/1999). What had been an orderly descent of thousands into the Bethnal Green tube station — part of a section of line that was under construction and being used as an air raid shelter — became a rapid descent as more than 1500 people surged forward to escape the sounds of war. This flight became a nightmare when falls on the steps occurred:

Either as a result of the sudden pressure from behind, or, by an unlucky coincidence simultaneously with the pressure reaching the people immediately behind her, a woman, said to have been holding or leading a child, fell on the third step from the bottom. This was observed both by a witness on the landing below and by at least 2 people in the crowd on the stairs behind her. As a result, or again, simultaneously, a man fell on her left. This occurred in the right hand half of the stairway. So great was the pressure from behind that those impeded by the bodies were forced down on top of them with their heads outwards and towards the landing. In a matter of seconds there was built an immovable and interlaced mass of bodies 5 or 6 or more deep against which the people above and on the stairs continued to be forced by the pressure from behind (Dunne, 1943/1999).

The real damage occurred as those outside the shelter — unaware of the problem on the stairs — pressed forward to gain the safety of the lower level. In the ensuing crush 173 people were killed (27 men, 84 women and 62 children) each one by asphyxiation.

Some details about the event were captured by the official inquiry (Dunne, 1943/1999): A Mrs. Barber, who witnessed the fall from behind, reported being lifted off her feet and carried down the stairs with her feet off the ground by the force of the crowd before the woman fell. The bodies could not be moved from below, and it took four hours to extricate the wounded and dead who were “pressed together into a tangled mass of such complexity that the work of extrication was interminably slow and laborious.” The conclusion of the inquiry was that there were several contributing factors to the disaster, including the lack of a ‘crush barrier’ at the top of the stairs that might have prevented the pressure of the desperate crowd from being directly applied to those on the stairs. Inquiry witnesses who had descended before the disaster were unanimous in saying that “the flow had been brisk and continuous” and that “while the people were hurrying and were talking of expected raid, there was no disorder.” Despite this, the inquiry report concludes with the suggestion that the disaster was caused by “a number of people *losing their self control*. . . and *practical common sense*, the display of which has hitherto prevented the people of this country being the victims of countless similar disasters” (emphasis ours); in short, the event was put down to panic — a point we shall consider in section 2.3.

This event illustrates two of the principles we discussed in the preceding sections. Notably, the speed increased when the relatively sedate crowd became frightened by the overhead explosions. This was a contributing factor to the fall on the stairs as well as the determined press behind the fall, and exemplifies Helbing’s ‘faster is slower’ paradox. Additionally, a failure of front-to-back communication occurred in this situation, as the crowd forcing into the station above the stairs was unaware of the unfolding disaster that their haste was perpetuating upon those at the front.

2.2.2 Lenin Stadium, Moscow

On October 20, 1982 an international soccer match was played between Spartak Moscow and the Dutch HFC Haarlem at Lenin stadium in Moscow. Some of the details of this event are different in official and unofficial accounts, although Reuters has referred to this event as the acknowledged worst soccer tragedy in the world. Fans began to exit the stadium down a dark

icy hallway before the end of the game.¹ When a late goal was scored, the sound of the cheering from within caused these fans to attempt re-entry. The re-entering fans met exiting ones — a single section of the stadium contained 10 000 fans — in what has been described as a ‘human mincer’. Although numbers of dead were initially reported as “little more than a dozen”, and eventually at 61 by Soviet authorities, an independent report seven years later put the number at 340, and this number is commonly used in references to the event for memorial purposes. (KLM, 2007; Ruizenaar, 2007; Crowd Dynamics, 2007; Edelman, 1993, p. 199)

Like the Bethnal Green tube disaster, this event highlights the failure of front-to-back communication as a large crowd — with no knowledge of an ongoing catastrophe — attempted to press forward as usual. Additionally, this event exemplifies Helbing’s ‘freezing by heating’ paradox as two crowds attempted to pass through one another with the high energy brought on by the excitement of an international sports match. It should be noted that the heterogeneity of the crowd goals was directly responsible for this disaster. Thus, a model wishing to study this disaster requires capability to model a heterogeneous crowd.

2.2.3 Who concert disaster

On December 3, 1979, in Cincinnati, Ohio, 11 people were crushed and died due to compressive asphyxia outside Riverfront Coliseum (now the US Bank Centre) before a concert featuring The Who. Interestingly, security personnel remarked after the fact that crowd conditions on December 3 were far from unusual for concert events, and that other concerts had been far worse in terms of crowding. This is a very well documented crowd incident, and police radio records and post-event interviews document the time-line fairly clearly. Two summaries of the disaster from original sources exist (Johnson, 1987b; Chertkoff & Kushigian, 2001), and it is these accounts which inform this section. As the communication problems in this disaster are so clear and well documented, it is worth considering this event in slightly greater detail.

The attendees began to arrive at the venue early in the afternoon for the 8:00 pm performance, gathering on a large terrace above ground level by designated entrances. This early

¹Whether fans were exiting of their own accord, or — as some versions have it — were forced to exit by police, is not important to our purpose here.

arrival was because 85% of the roughly 21 000 tickets sold were for 'festival' or general admission seating arrangements. The crowd grew to the point where individual control was lost, suggesting that — according to Fruin's quality of service norms — crowd density exceeded seven persons per square metre. In fact, a group of people at the rear were experimenting with the creation of shockwaves through the crowd by shoving forward in concert. One small 17 year old girl, who had been trying to escape the front of the crowd due to unbearable conditions, reported overwhelming crowd force almost an hour before the major incident:

I lost my footing an' slowly but surely began going down. People behind me could do nothing to stop the pushing. I was saying "No. No. Please help me..." Some of the people around didn't even hear me.... So then I grabbed someone's leg an' whoever that was told three other guys about me. They all pushed me up, pulled me up, but it was hard.... At about 7 o'clock I passed out. The four guys who pulled me off the ground helped me to stay up until we got through the door (qtd. in Johnson, 1987b).

By 7:20 pm one police officer estimated that 8000 people were in the crowd outside the doors. Entry was delayed due to the late arrival of the band and their need to do a sound check before the performance. At about 7:30 pm less than 16 exterior doors were opened into an intermediate ticket lobby and the crowd surged forward. About 25 people fell to the ground a short distance from the exit. Despite the efforts of those around them to assist (or at least avoid) the fallen, and to protect them from further assault, additional ranks of crowd members fell atop them or were forced over them; the pile grew to 3–5 people deep at its worst, was 10–12 feet in diameter, and some people were lying on concrete for as long as 30 minutes. Those farther back (just 10 feet back according to one interviewee) were unaware of the situation and continued to push to reach the stadium. After analysing numerous interviews, Johnson concluded that unlike those nearer the rear, crowd members nearer the front were surging forward to reach safety rather than to obtain good seats.

Security personnel tried to regulate the number of persons in the ticket lobby by closing the doors when it became full. Each time this occurred people in the vicinity of the exit became concerned that the venue was closing and the music from the sound check led them to believe that the concert was under way, adding to the crowd's determination to enter.

Police and security were unable to patrol the terrace and did not observe the pile of people. Due to their experience with rowdy concert-goers at previous events, and their belief that concert-goers were drunk or high (rather than injured), authorities did not grasp the magnitude of the disaster while it was unfolding. At 7:54 pm police became aware that a situation was developing, although they initially thought that the problem was a single victim of a heart attack. The extraordinarily high crowd density hampered rescue attempts and ultimately 11 died.

There were numerous types of communication breakdown in this crowd situation. Not only were many crowd members unaware of the unfolding disaster, many crowd members were unaware that people at the front were afraid for their lives well before the fatal pile-up occurred. Those toward the front of the crowd — even if unaware of the pile-up — were surging forward to safety, and those behind them were unaware of this reason for the surge and joined in thinking they were missing the concert. Concert and venue managers assumed that all was well, even when face to face with evidence of injured attendees in the ticket hall. Police apparently had no overhead view of the situation, no facility to patrol the crowd, no knowledge of the crowd conditions before the doors opened, no knowledge of the pile-up, incorrect information about the injuries when they did become known and no way to reach the injured. The only communication that functioned at all in the crowd was local communication amongst neighbouring crowd members. The quotation above reveals the kind of communication that did occur. Although Johnson did not study this directly, he states that local assistance was generally attempted, and that there was difficulty and frustration in communicating to the rear of the crowd.

From a modelling perspective, there is another interesting fact revealed by this event: not all members of a crowd act toward the same ends. In other words, this was a heterogeneous crowd. At first glance, this disaster may seem like a classic craze in which a crowd of essentially interchangeable people competes for a limited resource (i.e. early access to preferred seating locations). In this case, however, the crowd was not a monolithic block. Johnson reports that many people attempted to escape the front of the crowd even before the doors had opened, some in the back experimented or amused themselves with crowd shockwaves, others in the back remained pre-occupied with the obvious goal of seeing a concert they thought had already

begun. When we view the situation in this light it becomes natural to include situatedness, communication and information in crowd models to produce individuals that are heterogeneous, that is to say non-interchangeable, with memory and experience.

2.3 The myth of panic

The media is quick to label crowd disasters *panics*. At the time, for example, the Who concert disaster was widely described as a tragedy caused by a type of acquisitive panic called a *craze*. Johnson notes that the event was described as a “stampede,” fuelled by “mob psychology” that resulted in “seemingly selfish, ruthless behaviour” by participants. In the 2006 crowd disaster in Iraq (when many tried to exit an island over a small bridge due to concerns of a suicide bomb attack) the BBC reported that “almost 1000 people are known to have died in a stampede... Witnesses said panic spread over rumours of a suicide bombing.” (BBC, 2005) It suffices to scan newspaper reports of this type for many more examples. But was the type of behaviour observed at the Who concert disaster really a panic? What is panic, and what does such a label imply? Panic is defined as follows:

Panic (adj.) of the nature of, resulting from, or exhibiting a sudden wild, unreasoning, or excessive state of fear or alarm.

Panic (n.) A sudden feeling of alarm or fear of sufficient intensity or uncontrollable-ness as to lead to extravagant or wildly unthinking behaviour, such as that which may spread through a crowd of people; the state of experiencing such a feeling. Also: an instance or episode of such feeling; a scare. (*Oxford English Dictionary*, 2007)

Panic is defined as an state of unreasoning, unthinking or non-rational behaviour, but to the contrary most crowd disasters show the opposite kinds of thinking. This question about the use of the word panic may seem like an unimportant semantic distinction, but in fact it is very important to us; the value of modelling human behaviour in crowds turns on whether people in crowds can process information rationally and use information to make decisions about courses of action. In fact, it turns out that most situations are not well described by a view of participants as irrational.

Several researchers have specifically studied the question of panic in crowds. Sociologists Quarantelli and Dynes (1972) studied different kinds of disasters (some first-hand) over many years and reported few incidents of true panic, including amongst most of the people at the truly horrific Coconut Grove fire.² Analysis of events by Sime (1983) and Feinberg and Johnson (2001) have shown that family groups tend to remain bonded together during disaster and act to maximise chances of group survival rather than individual survival — contrary to what would be expected in an irrational panic. In his analysis of records and interviews from the Who concert tragedy, Johnson concludes that panic must include “unregulated competition” amongst crowd members; far from uncontrollable fear leading to unthinking competition Johnson found that almost every interview contained stories suggesting the opposite: intact social connections, people helping friends and strangers, people looking for shelter, people trying to communicate (Johnson, 1987b). Keating interviewed 100 people who escaped single-dwelling residential fires and found “no evidence as yet of any panic behavior by the people involved,” instead finding “just the opposite: altruistic, helping behavior — especially among family members” (Keating, 1982). Along the same lines, an official report determined that “September 11, 2001 showed that people are willing to help out others during an emergency, even if they do not know the person ahead of time” (Averill et al., 2005, p. 148). Nothing in the report suggests that widespread panic was found during the disaster; an author of the report believes that widespread panic did not occur amongst evacuees and noted that photos taken in stairwells during the evacuation show people smiling, interacting and waiting patiently (Guylène Proulx, pers. comm.).

Reviewers have come to much the same conclusion. Chertkoff and Kushigian (2001) state that their “case histories of good and poor traffic flow give little support to those theories that allot a major role to the occurrence of panic” as defined by high emotional arousal coupled with irrational behaviour. In a review of multiple events, Johnson “was struck not by the breakdown of social order but by its strength and persistence; not by the irrational, individual behaviour of popular myth, but by the socially structured, socially responsible and adaptive actions of those

²This fire on November 28, 1942 in Boston trapped 200 people in a basement room with one narrow staircase for escape. Competition was fierce for exit, but in the end many were killed by flashover on the stairs. Three hundred others were killed in other areas of this bar complex. The fire led to many modern fire code requirements like panic bars on exit doors, exit signage, mandatory sprinklers and multiple stairways for exit from basements. (Chertkoff & Kushigian, 2001)

affected” (Johnson, 1987a). Clarke notes that “after five decades studying scores of disasters such as floods, earthquakes and tornadoes, one of the strongest findings is that people rarely lose control” (2001). It seems that the view that people generally act irrationally in crowd disasters is not borne out by the evidence.

As we have seen with the Who concert disaster and at Bethnal Green, once we look closer at the situation there is a lot of rational behaviour to be found, even under extraordinary stress. At Bethnal Green, people’s flight from the sounds of war seem entirely reasonable. Pushing from the rear — whether intended to communicate an emerging social norm for urgency or to hurry the process — may well seem adaptive from the point of view of the individual at the back of the crowd who does not realise that people at the front are being injured. It is important to note that the emergent dynamics of multiple rational agents can itself be an irrational collective behaviour, as is seen, for example, in the tragedy of the commons (Hardin, 1968), although we shall have more to say about this below. One possible explanation of the appeal of the concept of panic is that people may find it easier to believe that the people in a crowd disaster must have misbehaved or ‘lost control’ than to accept that emergent effects like force and injury can be a consequence of rational — if stressed — behaviour. A non-adaptive emergent behaviour can arise from perfectly rational behaviour by individual actors.

For those within crowds we must also remember that Fruin tells us that individual control is lost. We must resist the uninformed judgements of the media, who, for example, condemned the “mob psychology” of “seemingly selfish, ruthless” people who “stomped 11 persons to death [after] having numbed their brains on weeds, chemicals and Southern Comfort” (qtd. in Johnson, 1987b) at the Who concert. In fact, those who were forced to stand on the victims did so with no control whatsoever of the placement or position of their body weight.

This is not to suggest that no situation has ever occurred in which people acted irrationally, or fled mindlessly in fear. However, researchers have not found enough of these cases to document this as a response pattern. In studying many crowd disasters in detail, however, they have found that viewing people as making rational decisions given the information available has more explanatory power.

2.4 Information in Crowds

The conclusion we draw from the results of the previous section is that people within crowds retain their ability to receive, understand, and process information — and indeed do so under even extreme circumstances. Thus, when aiming to study emergent crowd properties and crowd safety by modelling crowds at the microscopic level, some concept of information and communication should be included.

Communication and information is known to be a factor in many crowd incidents. Pauls' observation is not unique as inquiries into situations like the ones just described have often noted the problem that those only a short distance from injured people do not realise that anything is amiss (e.g. Johnson, 1987b). Proulx and Sime have noted that in fires notification is often delayed (sometimes for fear of causing a 'panic') resulting in precious minutes being wasted early in a situation. This delay can be self-fulfilling, with not enough time available for occupants to assimilate the information and move to exits once notification is given. Their experiments demonstrate that people do not panic in a perceived emergency, that they "act in accordance with the knowledge of a setting and situation," and that they require "information to define the situation and to take the decision to evacuate." (Proulx & Sime, 1991).

Fruin developed a concept that crowds are best modelled according to the acronym FIST (Force, Information, Space, Time) (1993). According to Fruin, the concept of information includes both pre-event factors (e.g. training) and intra-event factors (e.g. appropriate warnings, crowd knowledge about exits and conditions). While pre-event factors are beyond the scope of this thesis, intra-event factors can be modelled according to our methodology, and we take up this challenge in chapters 7 and 8.

In the last section we compared crowd disasters to a tragedy of the commons, suggesting that irrational collective behaviours can emerge from rational individual ones. The discussion of this section indicates a key difference between typical commons tragedies (e.g. environmental degradation) in which the speed of the emergent consequence (e.g. a polluted and dysfunctional ecosystem) is slow (relative to people's behaviours) and information is unconstrained. In crowds, people may not have the information to understand the consequences of their actions, and so

their behaviours may be rational *and not motivated egoistically*, while the emergent consequence remains catastrophic.

In section 9.1 we shall return to many of the points of this section, armed with the results of our demonstrations.

2.5 Summary

We have seen the essential principles of pedestrian movement and crowd dynamics: Pedestrians, unlike vehicles, have flexibility of movement and relatively unconstrained social norms for where and how they can carry it out. In crowds, people's movement becomes limited and effective levels of service decrease. Force in crowds can occur, generated by pushing and leaning, and tends to cause inefficiencies at bottlenecks, including characteristic rainbow-like arches around exits. Significant forces, which accumulate through multiple people pushing and leaning in concert, can reach levels where people are injured or killed through compressive asphyxia. Crowd forces lead to paradoxical emergent effects: slower collective progress results when individual speeds increase, and increased speeds reduce available time to accommodate groups that must infiltrate through one another. We reviewed three crowd disasters, and noted that panic is not a good explanation of crowd behaviour. Finally, we discussed the limited information flow in crowds, sometimes intentional on the part of facility managers, and sometimes a consequence of inadequate front-to-back communication.

These crowd basics will be referred to throughout the remainder of the document as we discuss existing models, and develop our demonstrations of the Microscopic Human Factors methodology. We now turn to a review of the literature in microscopic modelling.

Chapter 3

Previous models, critique and problem

Microscopic crowd models seek to understand crowd behaviour through simulation of individuals' behaviours. Through simulation of the individual dynamics in crowds, emergent patterns of crowd behaviour can be observed and correlated with known principles of crowd dynamics. This chapter begins with a literature review, which we have divided into two sections. First, we review relevant microscopic models, discussing their approach and their aims. Second, we present a critique of these models, from the point of view of Sime's concerns about the integration of physical and behavioural approaches to modelling, with a particular focus on the modelling of panic. In both the review and the critique, presentation of models is grouped thematically: first the microscopic crowd models themselves, second theoretical frameworks, and third the models with a specific focus on psychological relevance.

We have divided the review of the models from the critique for two reasons: First, the critique is made easier when we can discuss common aspects of models that have previously been introduced. Second, the critique marks a transition in this thesis, the point at which our focus shifts from the work of others and toward our own contributions.

We close this chapter by identifying general problems with the state of microscopic modelling: a lack of attention to Sime's concerns and assumptions of panic behaviours. This key section motivates the work of the rest of the thesis.

3.1 Review of previous models

3.1.1 Microscopic crowd models

Social Force model

Perhaps the best known of the microscopic crowd models is the Social Force model (Helbing, Farkas, & Vicsek, 2000b). The motivation for the model was research in social psychology, the conclusion of which was that mass panic was an appropriate model for crowd disasters; that, after LeBon (1895/1960), mob psychology homogenises behaviour across individuals; that interactions among people in crowds become primarily physical in nature; and that maladaptive behaviours like constant pushing lead to jamming at exits occurs due to un-coordinated movement, termed ‘incoordination’.

In the social force model agents aim to maintain a chosen velocity. Their actual velocity, however, is tempered by forces — some physical, and some the result of hypothesised psychological tendencies expressed as forces — that allow agents to keep their distance from other agents and walls. Each agent computes an acceleration dV/dt according to the following strategy (we have simplified Helbing’s formula for the present purposes):

$$m_i \frac{dV}{dt} = m_i \frac{v_0 - v_i}{\tau_i} + \sum_{\substack{j \in A \\ j \neq i}} [f_a(i, j) + f_c(i, j)] + \sum_{w \in W} [f_a(i, w) + f_c(i, w)] \quad (1)$$

This equation, for agent i , contains three summed terms. The first determines a corrective force that the agent exerts to move itself toward the desired velocity v_0 by multiplying the agent mass m_i by a corrective acceleration (difference between desired and current velocity v_i , divided by a time constant τ_i). The second computes the force that each other agent j (from the set A of all agents) exerts on agent i , composed of an avoidance force $f_a(i, j)$ (psychological) that mimics i ’s desire to avoid collision with j and, only for agents that touch, a contact force $f_c(i, j)$ (physical) that models body compression and friction on i due to j . The third term models avoidance and contact forces for each wall from the set W of walls in the model analogously to those between agents.

Each agent must therefore consider all other agents and the entire environment at all times, like an n -body problem in physics. Movement, however, appears quite realistic and the model is able to map physical body sizes and other physical measures directly into modelled space. Model results have demonstrated properties of crowds, including arching and jamming at exits, and the ‘faster is slower’ paradox.

4-Ped model

Blue and Adler developed the first cellular automaton (CA) simulation of pedestrian behaviour (e.g. 1998) and eventually extended this model to a four-directional pedestrian simulation on a toroidal grid (2000). Their modelling efforts were inspired by previous CA models of vehicle traffic, but they noted that the different conflict situations that pedestrians could be in required a pedestrian-specific model.

Blue and Adler were concerned primarily by highway traffic planning for pedestrians, and were not motivated by an interest in crowd dynamics *per se*. Their analysis was based on slowing of pedestrian traffic due to interference from crossing flows, initially in one dimension, and later with two, or m -dimensional flows.

Their model was able to replicate their observations of dynamic lane formation in Grand Central station in New York.

The Floor Field model

The floor field model (originating in Burstedde et al., 2001) aims to reproduce crowd behaviour in a cellular automaton, using principles analogous to those used by swarm intelligence (e.g. Bonabeau et al., 1991). As the floor field model will be used as the foundation of the methodology demonstrations of chapter 6–8, we shall confine ourselves in this review to the essential themes of this model; an in-depth description of the model will be presented as chapter 5.

The modelled space is divided into square grid cells, holding up to one agent per cell. The cells are also a convenient co-ordinate system for continuous or particle-based representations of information called fields. Just as a set of temperatures can be obtained from various locations

in a room, collectively referred to as a ‘temperature field,’ so can the grid squares anchor a set of information that has values at particular locations in the modelled space. This information can then be accessed locally by an agent in making movement decisions, formulated as a probabilistic selection of adjacent cells informed by a mathematical combination of the local field values.

The basic model contains two fields, one tracking distance from the exit (*static field*) and the other tracking agent movement over time in a manner analogous to ant pheromones (*dynamic field*). By balancing agents’ perception of these two fields, different kinds of crowd can be modelled. More attention to the static field causes agents to move quickly to exits. More attention to the dynamic field causes agents to follow one another in a pattern of motion described as ‘herding’.

In order to “avoid jamming transitions at unrealistically low densities,” the model contains an agent memory for progress. Agents become “unhappy” if repeatedly blocked and become “happy” again when movement resumes. This state transition corresponds to a move to a more random movement pattern and back again. This happy/unhappy concept was dropped in subsequent iterations of the floor field model.

This basic model has become the platform for experiments, extensions and other developments. The dynamics of the model were explored by Kirchner and Schadschneider (2002). An additional friction parameter was introduced by Kirchner, Nishinari, and Schadschneider (2003) to refine behaviour when there is contention for space. The model was later extended by Nishinari, Kirchner, Namazi, and Schadschneider (2004) to support arbitrary geometries rather than a square world. Extensions have been made to the model to support higher and varying velocities (Kretz & Schreckenberg, 2006; Kretz, 2007). The results have been compared to ant models (e.g. Nishinari, Kirchner, Namazi, & Schadschneider, 2006), and model results have been compared with evacuation data in an aircraft scenario (Kirchner, Kluepfel, Nishinari, Schadschneider, & Schreckenberg, 2003).

3.1.2 Modelling frameworks

The Mmass framework and SCA model

The multi-layered, multi-agent situated simulation model (Mmass) (Bandini, Manzoni, & Simone, 2002) is a framework that is more general than simple crowd modelling. The primary aim of the framework is to recognise that, as suggested by the principles of system science, situations unfold on multiple levels of complexity at the same time. The agent environment is described as a graph of available sites, allowing for arbitrary geometries. Agents are modelled as mobile or immobile and emit or consider specified fields, and change their receptive and transmissive states according to finite state machines.

In the Mmass model, an agent can participate simultaneously in two or more simulated multi-agent environments to find global maximum points of a measure of interest. For example, a model is developed that helps to locate promising sites for a shopping centre, with the shopping centre being an agent in two simultaneous layers. The first layer, a geographic one, is a rectangular grid corresponding to a physical city. The second layer, a logical one, is a completely connected graph with all shopping centres as nodes that models competition amongst centres. Through simulation, the framework allows both distance and competition functions to be optimised.

The Mmass framework has been refined specifically to study multiple pedestrians in a situated context. This refinement is termed the situated cellular automaton (SCA) model (Bandini, Federici, Manzoni, & Vizzari, 2006). SCA gives a formal specification of agent types and their possible states. Agents are sensitive to fields (broadcast by agents or objects) which guide motion or trigger state changes depending on their intensity and the state of the agent. The SCA framework has been used to model boarding, riding in and alighting from train cars. In this demonstration of the framework, agents can be boarding, standing (on board), seated (on board), detraining or exiting (the platform after detraining). Seats on the subway cars and standing positions are nodes in an irregular spatial graph, as are doorway areas and platform areas. The agent states determine the way in which agents choose to navigate the space, resulting in heterogeneous behaviour amongst agents in the same space (Bandini et al., 2006).

Co-Fields

The co-fields formalism is aimed at unifying multi-agent systems based on swarm intelligence under a single formalism (Mamei, Zambonelli, & Leonardi, 2002, 2004). Like the MMass framework, the co-fields approach is a general approach to problem solving that formalises a general category of solutions that researchers have been using to solve problems; in this case, the solutions involve fields (similar to the floor field model). The co-fields approach involves agents following gradients in a space of overlapping fields, some persistent, some generated dynamically by the environment and some generated dynamically by agents. The potential value of this kind of formalisation is that principles can be developed from general cases that are immediately applicable to more specific implementations of the ideas in the formalism.

The floor field model — even though it does not directly reference co-fields — is the major representative of the co-fields approach in crowd modelling. The co-fields formalism, however, has been used to investigate pedestrian situations (e.g. patrons moving through an art gallery looking for guides). Unlike the restricted CA implementation of the floor field model, the increased generality of the co-fields system allows agents to perform more processing than is anticipated by the floor field model. This is because agents are not restricted to a probabilistic movement system, are not forced to access information locally (although implementations may choose to do so) and can combine, process, and modify the fields in any way the modeller intends.

3.1.3 Models distinguished by their psychological focus

MACES, PMFserv and HiDAC

The MACES model (Pelechano, O'Brien, Silverman, & Badler, 2005) aims to consider what happens when a building is evacuated and some agents have limited knowledge of the building. It contains two separate systems. The first, or high-level, system is a wayfinding system that aims to allow communication of route knowledge between agents, allowing them to synchronise their mental maps. The second, or low-level, system is a social force panic-based evacuation simulation that exploits the high-level system's route finding behaviours.

An interesting aspect of the MACES model is the high-level wayfinding system. It divides agents into three categories: trained leaders, untrained leaders and followers. Trained leaders, such as firefighters, are presumed to have full knowledge of the building and will assist others. Followers will explore the building unless in the presence of a leader in which case they will simply follow. Untrained leaders will explore and lead followers, but do not have perfect knowledge of the building.

These researchers have noted the general lack of psychological relevance in crowd simulation, saying that current models “lack a wide variety of behaviours based on demographically and culturally validated behaviours” (Pelechano et al., 2005). In order to remedy this, they have decided to introduce a more complex psychological model for individuals as a complement to the MACES system, namely PMFserv.

The PMFserv environment used allows researchers to represent facts symbolically in long-term memory, allow short-term access to this data using a blackboard structure, and influence retrieval and processing of the information. Stress, coping styles, physiology and memory are variables in this system that determine access to cognitive resources. Available actions are rated based on an emotional model to determine a course of action, which is sent back to the MACES system for implementation.

The current conception of this model, HiDAC (Pelechano, Allbeck, & Badler, 2007), is aimed at realistic computer graphics of crowds, improves the efficiency of the model by confining processing of neighbouring agents to those within the current room, and also includes additional psychological effects including effects on the low-level system controlling pushing forces and desired velocities.

Fridman and Kaminka

Natalie Fridman and Gal Kaminka aim to find a single algorithm which would naturally account for many different crowd phenomena, based on the perceptions and actions of psychologically plausible individuals (Fridman & Kaminka, 2006). Their approach is motivated by LeBon’s view of crowd psychology (1895/1960) which suggests that individuals in crowds self-homogenise

through processes of imitation and contagion. Specifically, they have selected Festinger's Social Comparison Theory (Festinger, 1954) as a psychology for agents. According to this view, agents compare themselves with others — especially as similarities become higher — and take steps to reduce differences between the self and others. They have evaluated their model using lane formation, for which some empirical data is available, and have reproduced known properties of this phenomenon in corridors.

Virginia military simulation

Researchers at the Virginia Modelling Analysis and Simulation Center are working on simulating crowd psychology in the context of military models of troop strategy when simulated troops are deployed in urban environments. Accordingly, their strategy (Nguyen et al., 2005) is to model individual levels of aggression, but quickly move this to the level of groups within the crowd. The main psychological variable measured is the 'group aggression level' which is used to motivate and animate a group leader that the remaining individuals are assumed to mimic. Behaviours are tied to this aggression level such that an aggression level of 'avoidance' implies the behaviour 'flee from' and an aggression level of 'aggression lethal action' leads to the behaviour 'shoot at'. Their view of crowds is motivated by the suggestion that crowds are composed of sub-crowds and individuals (Musse & Thalmann, 2001). Although it is possible in this view to have group sizes of 1, the focus of creating sub-groups in the model is to reduce computational effort.

MASSEgress

The MASSEgress model (Pan, Han, & Law, 2005; Pan, 2006) is a microscopic crowd simulator that aims to study the effects of human behaviours such as competitive and queuing behaviours, understanding them from the point of view of known human behaviours in crowd environments. Like us, Pan was motivated by the body of work investigating human behaviour in crowd situations, and their lack of representation in crowd models. The research is aimed specifically toward understanding human behaviour in the context of building engineering and design. The

simulator emphasises a rule-based psychological simulation, providing the minimal physical and environmental simulation required to model agent movement within the space.

3.2 Critique of previous models

Have Sime's concerns been taken into account in microscopic modelling? We take Sime's concerns to be that crowd dynamics (and indeed all human movement) must be understood as an interaction between people and spaces. On this view, an appropriate model must take into account both psychological and physical aspects in simulating crowds. Although each of the approaches to modelling just presented has been instructive in its own way (indeed some have generated great understanding of crowd dynamics) we believe they do not properly take account of Sime's concerns. Let us consider each of the models described to substantiate this view.

3.2.1 Microscopic models

We have been guided thus far in our study by Sime's observations that crowd models in engineering do not take account of our knowledge of human behaviour. People in crowds retain their cognitive abilities and panic, as an irrational and self-defeating behaviour is so rare as to be undocumented on any significant scale in even the worst disasters studied. It therefore seems reasonable that microscopic models of crowd dynamics must avoid including irrational or panic-based behaviours on a large scale, and should instead put an emphasis on cognitive and behavioural simulation. In this, the microscopic crowd models we have considered are lacking.

Social force model

The social forces model is explicitly identified as a large-scale panic model, aiming to "model the collective phenomenon of escape panic" (Helbing et al., 2000b). This explains the general lack of deliberation by agents, and why they have homogeneous behaviour; its agents are interchangeable, operating on the same information to the same ends. This model is better viewed as a particle dynamics simulation, with the addition of a limited internal drive that allows

for inter-agent repulsion. The avoidance and social forces within the model, although suggestive of human factors, more closely resemble magnetic interactions than anything analogous to human psychology. This is similar to the attraction-repulsion framework used by some animal aggregate simulations (e.g. Boids: Reynolds, 1987) that has been criticised on the basis that it is a poor simulation of the mechanisms of aggregation, and leads to poor agreement with real aggregates (Schellinck, 2008). This, combined with the model's focus on unrealistic panic behaviours indicates that it does not take account of Sime's concerns.

Floor field model

The floor field model was inspired by a desire to extend vehicular cellular automata to the domain of pedestrian simulation (Burstedde et al., 2001). It has been used to study several different pedestrian phenomena — notably bottlenecks at an exit during evacuation (Kirchner & Schadschneider, 2002). The floor field model unfortunately follows the social force model in identifying panic as a likely outcome of some situations, but takes the positive step of modelling 'normal' conditions as well.

The normal conditions, however, do not seem to accord with Sime's view of what constitutes a proper understanding of human behaviour in crowds. The model specifically avoids modelling cognitive behaviours: "To keep the model simple, we strongly emphasise the principle to provide the particles with as little intelligence as possible and to achieve the formation of complex structures and collective effects by means of self-organisation. Effectively there is absolutely no intelligence... in our model" (Burstedde et al., 2001). In our view this is not a strength of the model; collective crowd phenomena must be understood as self-organisation, but this does not argue for non-consideration of human behaviour. Although doing so may simplify the model, it is our view that in so doing the model becomes oversimplified. It is our thesis that crowd dynamics must be viewed as self-organisation of humans, rather than self-organisation of particles.

The floor field model, as originally conceived, contains a parameter in each agent that controls whether it is ‘happy’ or ‘unhappy’. This labelling suggests an emotional model (suggesting a human factor) of frustration at being forced to remain still. The authors indicate it was introduced “to avoid a jamming transition due to obstacles at unrealistically low densities” (Burstedde et al., 2001) by reducing an agent’s commitment to a bad situation. Although it was introduced to solve a technical problem (rather than to specifically address the lack of human factors) this was a step in the right direction. This parameter was removed from the model in the next iteration, however, so it is fair to say that the floor field model is not intended to be a human factors model.

Due to its cellular automaton structure the floor field model considers all agents to be identical. Similar to the social force model’s avoidance forces, in the floor field model it is the following behaviour that seems less cognitive than magnetic. Unlike the social forces model the floor field model does not take account of crowd forces and so does not model agent safety. The floor field model, like the social force model, is an acknowledged particle simulator, rather than a simulator of humans in crowds. Accordingly we conclude that it does not take account of Sime’s concerns at present.

Despite these problems with the floor field model, it does make a very good base for a multi-agent system that can model human behaviours. This is because its discrete grid structure and local processing by agents is highly compatible with a multi-agent strategy and is quite extensible. Accordingly, while the model cannot take account of Sime’s concerns, we will adopt it as our basic model, using it to demonstrate (in chapters 6–8) how Sime’s concerns can be addressed within a microscopic model.

4-Ped model

Blue and Adler’s 4-Ped model is aimed at reproducing an observed movement behaviour of pedestrians, and does not aim to simulate independent pedestrians making individual decisions.

3.2.2 Modelling frameworks

Neither the Co-Fields or the Mmass models are explicit models of crowd dynamics. Both are theoretical frameworks that aim to document an approach applicable to the modelling of many different kinds of multi-agent and CA situations, respectively.

Co-Fields

The Co-Fields model suggests the paradigm that agents sensitive to distributed data structures may make local decisions that will result in certain emergent properties by the set of all agents. This research has been utilised to obtain certain conformations of agents in space and model typical swarms such as ants (Mamei et al., 2002) and to distribute humans in a gallery using radio-based network nodes that suggest destinations for tourists (Mamei et al., 2004). The latter study is less a study of crowd dynamics, and more a physical realisation of a computational problem.

Mmass and SCA

The Mmass framework, by contrast, is a formal specification of a general CA agent environment. Agents in the Mmass framework as currently conceived have a low cognitive capacity, as agents do not integrate multiple pieces of information found locally. Although Mmass agents do have a rudimentary internal state, changes in these states are triggered by simple presence of a single field above a threshold value.

It has been noted (Mamei et al., 2004) that Bandini's Mmass framework, although it does allow for rudimentary internal agent state, has an external trigger system that places the decision making outside the purview of the agent. The Co-Fields framework is more flexible in that agents are able to assimilate information and make decisions locally but it really represents more of a modelling guideline than an implemented system for modelling crowds.

Regardless of their individual merits or strengths, it should be noted that neither of these frameworks is a crowd model. In and of themselves, these two frameworks do not take Sime's concerns into account. Such account can in principle be taken by any modeller using the

framework to study crowds. The Co-Fields formalism has not been directly used to study crowds; the floor field model, discussed above, is the closest such implementation (although it was developed independently).

As discussed in section 3.1.2, the MMASS framework has been applied specifically to crowd environments in a formulation dubbed situated cellular automata (SCA) (Bandini et al., 2006). The SCA approach focuses attention on heterogeneity in crowds, allowing agents to change their movement goals based on extraction of information from the local environment. This is a positive step with reference to Sime's concerns, helping to understand people's behaviour in context, and the rigorous formality of the specification is an asset in understanding some of these behaviours.

We find that this formality, however, tends to make the approach rigid. Let us consider agent communication as an example: An agent a_1 in state s , broadcasts a field that, when perceived by another agent a_2 , alters a_2 's state. In reality people *decide* to communicate, and seldom communicate continuously (i.e. continuous yelling). The formal model can not accommodate this, however, as fields broadcast indiscriminately and continuously. As there is no room in the mathematical formalism to introduce logical operators that might simulate selective or probabilistic communication, interruption of the broadcast would entail altering the agent's state. Creating additional states to account for logical processing like this could overwhelm the model (or modeller), and introduce questions of how these changes of state would be triggered within the formalism.

Our view is that describing human factors in a computational model rather than a mathematical framework affords more flexibility to create human behaviour simulations at the microscopic level. Additionally, the SCA model does not provide for a force or safety simulation. In this respect it may depart too much from the physical situation being modelled and is limited in terms of the effects that can be studied. Although the SCA framework is on the right track in terms of increasing relevance to people, we will adopt a different approach.

3.2.3 Models distinguished by their psychological focus

Fridman and Kaminka

Fridman and Kaminka are approaching this same divide between engineering and social psychology but start from the opposite side. They have looked to psychology in an effort to find a single theory that can explain human aspects of crowd behaviour, providing different crowd behaviours in different contexts. They have adopted LeBon's explanation for 'mob psychology' in which people in groups tend to compare themselves and minimise difference (1895/1960). The authors argue from LeBon's position that an essential aspect of crowds is their homogeneous nature: "the participants act as if governed by a single mind, despite using little or no verbal communications" (Fridman & Kaminka, 2007). Accordingly, they have adopted Festinger's social comparison theory (SCT), taking a microscopic view and aiming for "a single computational mechanism — a single algorithm — which would account for different crowd phenomena by virtue of the actions and perceptions available to each individual" (Fridman & Kaminka, 2007). We do not believe that the single mechanism of Festinger's difference-reduction theory can account for human behaviour in crowds, and so on this point we disagree with the authors. The position that people in crowds do not communicate, behave identically and homogenise is in direct conflict with Sime's observations and our discussion of human behaviour in crowds (chapter 2), particularly Johnson's reports that crowds are anything but homogeneous. In our view the model risks oversimplifying cognition in crowd behaviour.

One reason to focus solely on a single psychological theory (such as SCT) to introduce human behaviours into a microscopic crowd modelling context is to evaluate whether that theory's predictions are sufficient to explain mechanisms of social behaviour within the individual. Occam's razor argues that such a theory, should it be shown sufficient, would be superior to a looser collection of interacting theories of internal psychological forces. Our interest is slightly different, however, as we make no claims as to the underlying mechanisms for social psychology, aiming instead to model emergent crowd-level consequences of social behaviours that are known (or suspected) to be relevant to crowd dynamics. Thus for us what is important is the ability to represent these behaviours at the heart of the microscopic model; the flexibility to use multiple

computational techniques is an asset, and so we also reject the contention that a model that uses a single mechanism (SCT or any other) is a better computational model. We prefer a strategy that more loosely allows for different cognitive behaviours to be simplified differently, and implemented in the best way that allows them to be fused with the microscopic rule set.

Despite our disagreement with Fridman and Kaminka on the points of SCT as a single psychological and computational model for crowd behaviour, there is much to commend in the approach taken. In particular, the authors provide a microscopic implementation of SCT, using it to analyse group formation and lane-formation in bidirectional flow. Microscopic implementation of any human behaviour in crowd dynamics is welcomed in any form. Also, we do not wish to claim that difference reduction *never* occurs in crowds. A more likely conception of this phenomenon may be localised within the crowd (e.g. a crowd member helping injured people after seeing another doing so). There must, however, be a limiting factor for this as we have seen that crowds are heterogeneous.

Virginia military simulation

This military model considers emotional-cognitive behaviours but the focus of the model is to generate behaviour for battlefield simulations, not to create crowd safety simulations. As such, the emphasis on simulating aggression seems like a mismatch for the questions we wish to answer. Indeed, incidents of aggression are far from the forefront in crowd safety research - researchers report widespread helping behaviours and miscommunication instead. The model bonds a psychological simulator to a separate movement simulator, an approach we discuss below, and which we do not favour.

MACES + PMFServ

Like several models and frameworks discussed above, the MACES + PMFServ model takes an interest in human behaviour in crowd modelling (Pelechano et al., 2005), in this case in a building evacuation scenario. The model explores the contribution to behaviour that occupant roles (e.g. leader vs. follower, trained vs. untrained) play in evacuations. Exploring social roles

in the context of a crowd model is applicable to Sime's concerns, and we shall mention this in section 7.7, although we shall be leaving it to future work.

Despite the promising research goals, however, we find the model is not conducive to our present purposes for three reasons. First, it uses the social forces model as the basis of its crowd movement simulation. This leads to a focus on panic behaviours and movements that are out of sync with our current conception of crowd behaviour, and seems at odds with the researchers' statement that models are lacking in "demographically and culturally validated behaviours" (Pelechano et al., 2005).

Second, the model is not set up for microscopic communication. Agents within a space are assumed to be able to communicate over any distance. Although this may be appropriate in tiny compartments such as in an apartment, it is not appropriate for spaces of larger size, as we have seen in Johnson's reports of the difficulty of communication at Riverfront Coliseum (section 2.2.3; Johnson, 1987b). More realism in perceptual mechanisms is desirable for microscopic crowd modelling.

Third, the modellers have chosen to couple the panic movement simulator to a large cognitive architecture (PMFserv). This architecture considers emotion, stress, and environmental factors on information retrieval and processing capacity, and includes a model of short-term memory for both spatial information and declarative facts. While this complex cognitive psychology is not out of place in a crowd simulation it may be so in a microscopic crowd simulation; the quick step from no psychology to a complete cognitive architecture is a very big step to take. Observation of a macroscopic psychology is difficult, and this precludes a connection between rules and emergent behaviour. We aim to identify the effects of introducing human behaviour into the model without raising new questions about what part of the cognitive architecture is responsible for the new emergent effects. See section 4.2.1 for more on this point.

MASSEgress

The MASSEgress framework (Pan et al., 2005; Pan, 2006) is a model aimed specifically at engineering buildings for safe egress. Like us, the authors express concern with the lack of

human behaviour in pedestrian modelling. This concern leads to the creation of a new model, a cognitive simulation that is based on decision trees. Our approach differs; following Sime, we believe that cognitive factors must be balanced by physical and environmental factors in crowd simulation. The MASSEgress framework bases crowd behaviour largely upon psychological and social dynamics, with a simplistic physical simulation and no force modelling. This limits its applicability in situations where large numbers of people are interacting and physical forces interact with social and psychological factors.

The MASSEgress model was almost entirely demonstrated as an egress simulator. Despite the focus on human behaviour, the MASSEgress research appears to have missed the opportunity to study how the introduction of these behaviours into a crowd simulation affects outcomes (and by extension, how the absence of these behaviours may represent a problem with current models). There is minimal investigation of how variation in human behaviour affects egress times, and no ability to simulate how changes in human behaviour affect pedestrian safety. The focus seems to be upon using the model to predict evacuation times rather than to understand emergent crowd behaviour.

3.3 Discussion: The problem with microscopic models

We are not aware of a microscopic crowd model that takes account of Sime's concerns by integrating a simulation of microscopic physical factors (such as crowd pressures and forces) with a microscopic behavioural simulation. We have considered a number of microscopic models of pedestrian movement and crowd dynamics. Some are concerned predominantly with physical and movement simulation; others focus on psychological and behavioural factors. No model reviewed has satisfactorily integrated these two aspects of crowd behaviour in a microscopic context. This means that the microscopic modelling community is not able to study emergent crowd behaviours in the context of individual rules based on both physical and behavioural factors. As real-life crowd situations involve interaction of both of these domains, it is our contention that the lack of such a model is an impediment to understanding crowd dynamics.

As we have seen, panic is a compelling concept that continues to crop up in crowd models — particularly in physics-based models, but also in some behavioural models. Although the problem may be lack of knowledge of the social sciences' arguments against panic, there is another reason why panic is compelling for crowd modellers: Because panicked individuals are assumed to exhibit behaviour that is irrational and free of social constraints, the modeller is relieved of the responsibility to simulate cognitive/psychological/behavioural aspects of crowd dynamics. The result is agents that are non-cognitive; specifically they are non-rational (taking actions that do not make sense), non-deliberative (containing no mental state or processing) and interchangeable. This in turn reduces the complexity of the model, and encourages techniques used traditionally to model inanimate objects. Although the panic assumption may be convenient in that it reduces complexity, it means that the behaviour modelled is at odds with our concept of what people do in crowds. This opens a model to the criticism that its results cannot be meaningfully applied to people because it is not a model of people.

3.4 Summary

We have reviewed a number of microscopic models and found them to be unsatisfactory in taking account of Sime's concerns. It is our contention, then, that the lack of a panic-free model that microscopically integrates both physical and behavioural factors is an impediment to understanding crowd behaviour, because only such a model could be legitimately applied to crowds of people. Our identification of this problem, then, leads directly to our research objective.

Chapter 4

Microscopic Human Factors

In this chapter we begin in earnest the original work of this thesis. In chapter 1 we introduced Sime's comments on balancing behavioural and physical conceptions of crowd dynamics. In chapter 2 we discussed (among other topics) the importance of viewing people in crowds as non-panicking rational actors, acting on the information at hand. Yet we have seen, in chapter 3, that Sime's comments are still very relevant in the case of microscopic models.

Inspired by the problems with the prior literature discussed in section 3.3, we first formalise and argue for the importance of our research objective, which is to devise a method for solving these problems. After determining what is missing from current models (relevant human behaviour), we then fully describe our proposed solution: the Microscopic Human Factors methodology. We present an argument for its utility and the relation between models created with the methodology and the world. Future chapters will provide demonstrations of this methodology in action.

4.1 Research objective

In formalising our research objective, the main problem identified by the literature review bears repeating: The lack of a panic-free model that addresses Sime's concerns is an impediment to understanding crowd behaviour. The central focus of this thesis is to argue that the integration of human and physical factors in microscopic models is necessary, possible and practical. The research objective of this thesis, then, is specifically as follows:

Research objective: *To develop a methodology to integrate relevant human behaviours with physical factors at the heart of microscopic crowd models for the purpose of improving relevance to people and studying emergent crowd-level effects.*

This objective has value and importance in the field of microscopic modelling. This is demonstrated first as the objective raises a number of interesting topics and questions that will be addressed in this and future chapters, including:

- what human behaviours are relevant?
- how should such a methodology be formulated?
- how does such a methodology guide model development?
- what validity and utility do models built with the methodology have?
- does the methodology work in practice?
- can the methodology help us model behaviours that are not well understood?

A second way that we can see the value in this objective is with respect to Sime's concerns. He has argued that we need models that integrate behavioural and physical approaches to crowd dynamics. Microscopic models that do take this approach will be an improvement over previous ones (according to Sime's concerns). We saw in sections 1.1.3 and 3.3 that part of the problem may relate to the difficulty of generating models that account for human behaviour. The lack of satisfactory models, as discussed in chapter 3, may be further testament to this difficulty. A project to develop a methodology that assists microscopic modellers with this process therefore makes a contribution.

A third factor that argues for the utility of the objective is that our models, which are outputs of this methodology, have met with interest from microscopic modellers. All three of the demonstrations of the methodology that we will present (chapters 6–8), as well as a short paper concerning the methodology itself, have been accepted on the basis of peer-review for presentation at scientific meetings or, in one case, in a journal (Henein & White, 2005, 2006, 2007, in press-a, in press-b).

The latter point argues for the importance of the methodology, but also points out that the methodology itself is not the only contribution made by this thesis. The demonstrations are contributions in their own right; as we shall see, each has shown that the methodology is practical and generative by producing interesting results, discussion, conclusions and suggestions for future work. When we review the conclusions and contributions of this thesis in chapter 10, these further contributions will be included and itemised.

With the objective described, and its importance asserted, we now generally turn to the work of creating the methodology. In the next section we take up the first question from our list, above: What is a relevant human behaviour? Answering this question will carry us to the specifics of the methodology.

4.2 Relevant human behaviour

As stated in the research objective, our methodology targets the integration of *relevant* human factors into microscopic models. We cannot model everything about people, and so we must narrow the scope of the methodology to what is most important.

4.2.1 The relevance of information processing

What are the most relevant human factors that must be incorporated? Sime's ball-bearing criticisms were directed toward treatment of people as non-thinking objects and we generally agree with this focus on cognition as the defining characteristic of people in crowds. (Panic, which could have derailed this position through an argument that cognition is essentially disabled in crowds, has been discredited as non-explanatory.) Dawson has asserted that the hallmark of cognition is information processing (1998), and the central assumption of cognitive science is that "the human mind is a complex system that receives, stores, retrieves, transforms and transmits information" (Stillings et al., 1987). This suggests introducing cognition into modelled agents through information processing.

Accordingly, one way to approach the problem would be to simulate a complete cognition for agents in the form of a cognitive architecture (e.g. ACT-R, Anderson & Lebiere, 1998).

Although this approach might improve relevance to people, it does not satisfy microscopic modellers because it negates one of the key benefits of microscopic models: the establishment of causal connections between emergent macro-level crowd effects and their origins in individual behaviour. Simple agent rules give way to a complex cognitive simulation that tends to be analytically opaque (as cognitive architectures are emergent systems themselves). This can lead to problems of scale in which the macroscopic cognitive layer overpowers the microscopic crowd model; the latter becomes little more than an arena, relegated to simulation of the physics and movement of agents, and is divorced from the real action unfolding (potentially on a different time-scale) within the cognitive simulator. This disconnection obscures the causal chain and it becomes difficult to explain the macro-level crowd effects by reference to processes within the individual.

Although cognitive architectures have many benefits, incorporating them at the heart of a microscopic model would be a challenge. It is not necessary, however, to import a complete cognitive architecture in order to model human behaviour in microscopic models.

4.2.2 Behavioural repertoire

Any model is, by definition, an abstraction that discards detail through simplification in favour of explanatory power. Because of the focus on emergent dynamics in microscopic modelling, current models abstract and simplify interactions in the domain of the individual; they do so by emphasising effects rather than mechanisms at the individual level. In the floor field model, for example, representation of physical space is highly stylised and so is knowledge of other agents. This is what allows a microscopic model to use simple rules at the individual level to produce complex emergent effects at the crowd level. In order to produce a scale match, our integration of human behaviour must follow similar design criteria.

With this in mind, when applying Sime's concerns to microscopic models, we accept that cognition-free models are oversimplified, but it is not *only* cognition that is lacking from the models: Human behaviour implies not only a mechanism for processing information but also a rich behavioural repertoire available for deployment by this cognitive mechanism. When we call

microscopic models into question because they ‘do not do what people do,’ we are bemoaning the lack of this behavioural repertoire.

As microscopic modellers of behaviour, we are less concerned with processes internal to an individual and more concerned with the behavioural *actions* of individuals — as it is these that collectively produce emergent effects. Our methodology for integrating human behaviours, then, recognises the microscopic focus on modelling effects rather than mechanisms at the individual level. In a microscopic context, cognition is important to behavioural simulation *in its role as support for behavioural repertoire*; it is, then, secondary to behavioural repertoire. In other words, it is the absence of repertoire — more than the specific cognitive process that underlies it — that calls into question the models’ applicability to human crowds.

On this view, we can take Sime’s ball-bearing complaints into account by extending the behavioural repertoire of agents. Our approach is to implement cognitive behaviour and behavioural actions at the same level of abstraction at which the microscopic models currently simulate movement (rather than representing formal models of cognition through a cognitive architecture). Consequently, our methodology expresses human characteristics as simple local rules,¹ fusing them with the simple local rules of the microscopic model. We call these simple local rules microscopic human factors (MHFs). The methodology is characterised by simplified information processing that triggers behaviour and supports its unfolding over time where required. This simplified view of the human behavioural repertoire reduces the complexity of the cognitive simulation required; this brings the human factors modelling into scale with the existing aspects of the microscopic model. This technique ensures that the causal connection between emergent crowd behaviours and the local agent rules is not broken.

In short, by reducing the fidelity of the cognitive simulation (instead focusing on behaviour) we adopt the same strategy that microscopic modellers have used to simulate physical aspects of the simulation. This reduced fidelity is warranted because our goal is not to explain intra-agent processing but rather to induce emergent effects by grouping many agents together.

¹A terminological note is in order on the use of the word *rule* in describing mechanisms of microscopic models. We use this word in the broad sense of “a principle, regulation, or maxim governing individual conduct” (*Oxford English Dictionary*, 1989).

Let us take the floor field model as an example of the value of including information-driven behaviours. Currently in the floor field model, interchangeable agents have one fixed movement goal that is not changed based on information obtained from the environment. This limited behavioural simulation is at odds with both Sime's concerns and also with our everyday experience of moving through crowded spaces. We know that people do not all have the same goals, and, especially in unfamiliar environments, may change goals multiple times. We can alter the microscopic simulation by providing a simulation of changing goals (see chapter 7). This alteration, even if based on a simple information-processing framework within an individual agent, can produce an improvement in our understanding of collective crowd behaviours as per Sime's arguments.

We now turn to an explanation of the Microscopic Human Factors methodology for adding a human behaviour to a microscopic model.

4.3 The Microscopic Human Factors methodology

The purpose of the methodology is to provide a framework for describing human behaviours that guides modelling of human behaviours as simple rules, integrated tightly with microscopic models. The methodology is general; while we demonstrate its use concretely in the floor field model (chapters 6–8), it is equally applicable to other microscopic models as implementation formalisms.² The methodology does not place limitations on the types of behaviours that can be modelled, but rather enforces the perspective of the individual in the crowd in explaining crowd mechanisms.

Our depiction of the methodology begins from a description of the behaviour in real crowds, guiding the modeller to focus on crowds as crowds, and not on other systems (e.g. seeds in a hopper) as metaphors for a crowd. We then explain how the methodology makes the theoretical contribution of the model explicit, by specifying the essential aspects of the behaviour in the language of individual people. We explain the importance of selecting a good formalism for

²Although the methodology is neutral on the question of implementation formalisms, that does not mean that all such formalisms are equally well suited to microscopic modelling of human factors. Our preference is for multi-agent systems, and we have discussed the appropriateness of this formalism in sections 1.3.2 and 5.1.

translation of these essential aspects into an implemented model. We shall then turn to a discussion of the validity of models built using the methodology.

4.3.1 The three levels of the methodology

The methodology requires three distinct descriptions of a behaviour of interest at three different levels (see figure 2). Each level, or layer, is a complete description of the behaviour at a different level of abstraction. Each of these three distinct descriptions is an important part of the methodology, and each level provides a locus for a different conversation about the crowd behaviour, and its origins in individual behaviour. By separating the description into levels, these conversations are made more precise, and each offers an opportunity to discuss a different aspect of the modelling process.

We shall consider each level, its purpose in the methodology, and the types of questions that can be asked at each level. As we move through the levels we shall use a trivial example human behaviour to illustrate the three descriptions: that of selecting a line at a supermarket (figure 3).

Specification. The first level of description is a clear statement of how a crowd behaviour, or aspect of crowd dynamics, occurs in the world. This may be a description of a well-understood behaviour, or a hypothesised behaviour for study. Its purpose is to develop a theoretical *specification* or mechanism for the behaviour proposed. As real cognitive systems — unlike microscopic cognitive simulations — are highly interconnected and complex, no model can account for all of crowd behaviour in one simulation. It is important to be clear about what is being modelled, and this level places a boundary around the simulation of interest. A clear specification of what is being simulated provides both a starting point for modelling, and a desired end point for microscopic simulation. At the specification level, discussion is at the level of real-world crowd behaviour and there is no need to discuss agents or modelling practicalities. Disagreement at this level suggests uncertainty about how people behave in crowds and why; in the case of more speculative research it is at this level that we present our hypotheses about crowds.

In figure 3 the specification is of the hypothetical type. Note that the description is given in terms of the behaviour of real shoppers. There is no mention of agents, nor of modelling

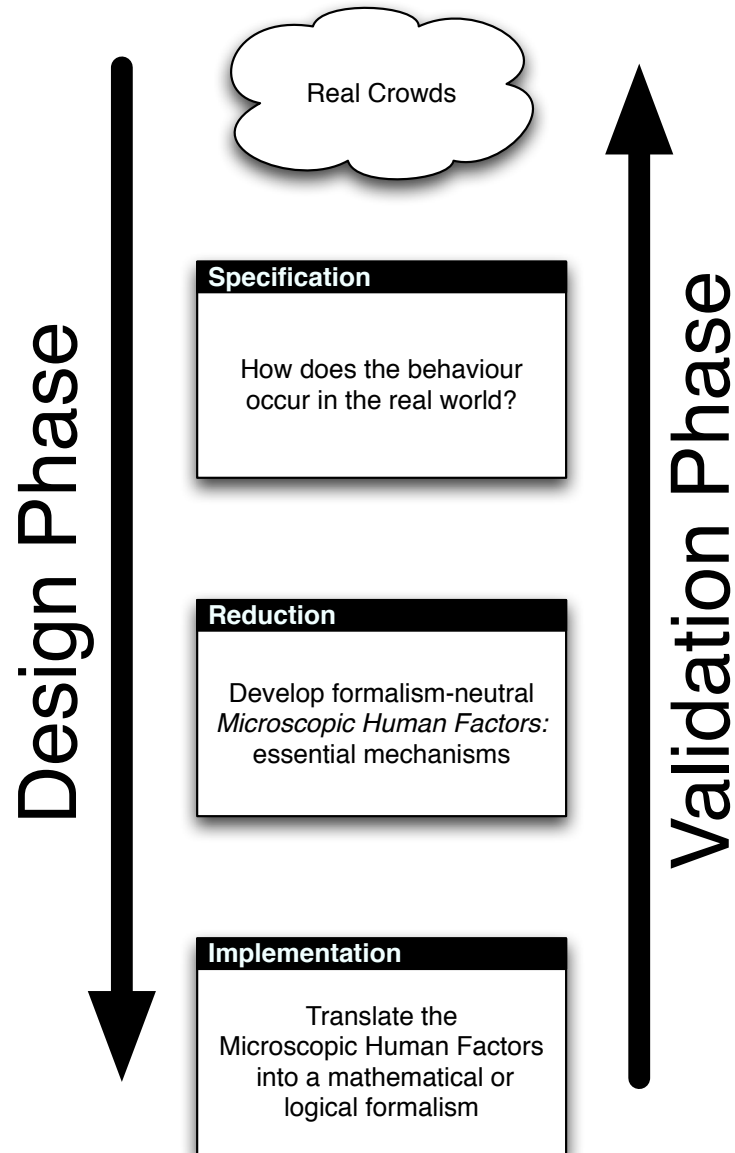


Figure 2: The three layers of the *Microscopic Human Factors* methodology. Each layer is a complete description of the behaviour, inspired by knowledge or hypotheses about real crowds. During the design phase the layers are created in order of increasing abstraction and specificity. During the validation phase the validation of underlying layers supports the layers above.

Specification. The specification requires a description of the way the behaviour works in the real world. We hypothesise that people make several types of observations and judgements when selecting a checkout line, in the interest of spending the least time in line. People evaluate the number of shoppers, the contents of others' carts and make judgements about the probable efficiency of the cashiers. We suppose that people prefer the express line to the regular line.

Reduction. Although there are several interesting phenomena at work in the specification, we must extract the essentials of this behaviour, and present it in an individual-focused, formalism-neutral manner for modelling. Although there are several judgements and decisions being made in the specification, we have chosen to reduce these to time as the essential prediction made by individuals. In other words, our model of line selection reduces, in a microscopic context, to an individual's time predictions. Our reduction assigns an approximation of the total number of items (across carts) in a line to be the key predictor of time. It is reasonable, however, to assume that the number of carts will also be considered, as will the type of line (express or not); these factors will be modelled as increases or decreases in each line's time estimation, respectively. The judgements about the cashiers, although interesting, are determined to be non-essential in a person's judgement of time, and are left to future work.

Implementation. At the implementation level we introduce the notation and formalism appropriate to the microscopic model of interest. In this trivial example let us simply assume that the model provides a set of lines $L = \{l_0, l_1, l_2, \dots, l_n\}$, a determination, $e(a)$ of whether each line is an express line (with $e(a) = 1$ if line l_a is an express line, else 0), a set of carts $C_a = \{c_{a_1}, c_{a_2}, c_{a_3}, \dots, c_{a_m}\}$ in each line l_a , and a weight of each cart C_{a_b} given by $w(a, b)$.

The implementation of our line selection algorithm is as follows. First, each line is assigned a unitless time prediction t_a (shown below) then the agent selects the line with the lowest predicted time. The score is determined by summing the weight of the carts in line, penalising each cart with 10 units of weight, and dividing the weight by 2 for express lines.

$$t_a = \frac{1}{1 + e(a)} \sum_{b < |C_a|} (10 + w(a, b))$$

Figure 3: A trivial example of the Microscopic Human Factors methodology at work. The task is to model selection of a line at the supermarket.

formalisms or constraints. We simply describe the behaviour that is to be implemented and how we believe it works.

Reduction. The second level of description involves *reduction* of the behaviour to simple local rules, and it is these local rules that we term *microscopic human factors*. This stage is the most important step in the methodology, fulfilling both a theoretical purpose and a practical one.

It is at this step that an abstraction of the behaviour is developed, a theoretical position that describes the *essential* aspects of a microscopic mechanism to underlie the behaviour. The essential aspects detail what information is used by individuals to trigger or carry out a behaviour, local sources of this information and how it is processed. They also cover physical constraints, situational factors and physical abilities of the individuals and how the behaviour is expressed over time. In short, this level requires a complete mechanism for the behaviour. Because microscopic models argue from individual rules, the description at this level is from the perspective of the individual within the crowd. Determining these essential aspects amounts to a theoretical position about a mechanism sufficient to reproduce the behaviour in human terms. Disagreement at this level suggests disagreement about how a particular crowd behaviour occurs as an interaction between people and space.

From a practical perspective, understanding of the reduced description does not require any modelling expertise because the description of the behaviour at the reduction level does not address implementation issues (which are deferred to the next level). This is what we mean when we say that the reduction is formalism-neutral. As with the specification level, discussion centres on people, their capabilities and behaviours, and the effects of the situation upon them. By keeping the discussion at the level of people while developing an abstract simple model that captures the essentials of the specification, non-technical researchers with considerable domain expertise can participate in the conversation surrounding essential elements of crowd situations and behaviour. This is a key factor in taking advantage of interdisciplinary co-operation to increase our understanding of crowd behaviour.

The reduction level exists, in part, to require a separation between simplifications introduced for parsimonious understanding of crowd behaviour and simplifications introduced for reasons

of modelling practicality. The reduction must be a defensible abstraction from a behavioural perspective in order to ensure that microscopic results are ultimately applicable to human crowds. It is not that modelling simplifications are inadmissible, but they are not hypotheses about crowds and must be introduced for what they are: implementation constraints. By making the reduction level explicit, the implementation formalism that most naturally captures the microscopic human factors can be preferred.

The points of the preceding paragraphs are demonstrated in the trivial example of figure 3. First, although the reduction is a complete description of the behaviour, we have performed an abstraction, focusing on *time* as the essential representation made of the lines by the agent. This is a theoretical position as to how the line selection behaviour essentially works. Consequently, we provide for a correspondence between the information available to the person (load in the carts, number of carts, line type) and the estimated time. We have identified an aspect of the specification (cashier judgements) as interesting, but non-essential. Second, the reduction is formalism-neutral and individual-centred; this means that the reduction has not pre-supposed what kind of model should be used, but simply focuses on what we believe to be true from the individual shopper's perspective. (The reduction is equally well suited, for example, to implementation in a numerical simulation or in a multi-agent system, but only from an individual's point of view.) The reduction shows what we think is important about the specification, without confusing matters by what it is practical to model using any given approach. Third, the reduction keeps the discussion about the essential mechanisms in crowds accessible to domain experts in supermarket planning, even those who may not be familiar with the formalism selected for the implementation.

Implementation. What remains after the two levels just described is the level of *implementation*. The purpose of the third stage is to harmonise the reduction, a functional description of the behaviour from an individual perspective, with an implemented microscopic model. This level involves a translation of the rules into a specific mathematical/logical formalism provided by a microscopic model. In the case of a new model, this may involve also the selection of an adequate formalism. It is at this stage that the behaviours must be implemented in such a way

as to interface closely with the existing rules of the microscopic model. An emphasis at this level on scale match results in intertwining physical and behavioural rules, and the preservation of the model's ability to connect individual action to emergent effects. Disagreement at this level suggests the mathematical rules or model logic are not a true translation of the reduced rules into the model.

This level brings with it not only the advantage that, once implemented, the behaviour can be simulated, but also that the requirement to implement is itself a test of the completeness of the specification and reduction.

In our supermarket example (figure 3) we do not go so far as to implement and test the behaviour (see chapters 6–8 for full demonstrations). There are, however, a few points worth mentioning. First, the line selection is done directly using the information represented in the model concerning lines, carts, etc. We do not need to import a separate line-selecting cognitive architecture to sit alongside the existing model. Second, the model did not provide for an exact translation of the reduction into the implementation. That is because the reduction expresses time in terms of number of items in carts, while the implementation works in terms of cart weights. Assuming this is a significant distinction, the modeller would need to argue for the appropriateness of this alteration in the behaviour (perhaps by drawing a parallel between weights and item counts in grocery store cart studies).

4.3.2 Validity and utility of the methodology

Validation refers to the process of determining the relationship between a model and the world. In this section we explore the extent to which microscopic models can be validated, what they can tell us about the world, and how the MHF methodology increases the validity of microscopic models.

Validation of microscopic models

One thing we should not expect of the MHF methodology is that it will suddenly allow us to do what microscopic models generally cannot: make precise statements concerning specific

cases/situations in crowd dynamics. Microscopic models are best at studying crowd behaviour in abstract situations, providing qualitative evidence that collective behaviours arise from specific individual behaviours, and demonstrating mechanisms that can account for these collective behaviours. There are several reasons why direct validation of microscopic models with human subjects is difficult:

First, many emergent crowd behaviours of interest test the limits of a physical setting to accommodate large numbers of people. This limits the potential for direct experimentation due to safety concerns for participants.

Second, as microscopic modellers we are interested in understanding interactions of specific behaviours in a massively complex system involving large numbers of cognitive actors with incomplete information and no central co-ordination. This is very different from the types of hypotheses that may be equally interesting and are testable in a simpler scenario, such as a laboratory study to track the effect of a clear independent variable (e.g. signal type) on a simple dependent variable (e.g. reaction time). The category difference here occurs because it is precisely the complexity of the situation that is of interest. The microscopic approach proceeds specifically through abstraction of the behaviours and a focus on qualitative evaluation. Even in a laboratory environment, people cannot be sufficiently controlled to isolate the effect of a behaviour of interest in the manner of a microscopic implementation. Direct human experimentation, in short, is not a match for the qualitative, abstract focus of the microscopic model that aims to draw general conclusions about behavioural interactions of agents.

The fact that testing with human subjects is not a good match for our purposes does not mean that the model has no validity. We take validity to mean relevance to the model's prototype: human crowds. Microscopic models have much to offer in the department of relevance to real crowds, and as we shall now see, the MHF methodology supports these arguments of validity.

Validity of the MHF methodology

One important benefit to the three-level methodology proposed is that, by isolating three different views of the model, multiple points are introduced at which the validity of the resultant model

can be naturally discussed. In section 4.3.1 we presented the levels in order of increasing abstractness. Once the simulation has been made, however, the order of levels can be reversed for the purpose of considering model validity. As we move upwards through the three layers we can use an implemented model that has adopted the Microscopic Human Factors methodology to answer different validation questions.

Implementation. At the level of the implementation we determine whether the microscopic human factors have been appropriately implemented. This includes simple checking, through experimentation, that the implemented model embodies the microscopic human factors described at the reduction level. It also includes determining that an appropriate implementation formalism was suggested that does not deviate from the reduction by introducing excessive simplifications for reasons of implementation constraint. (Certain reductions may be very difficult to implement, for example, in a purely numerical simulation. This is because cognitive factors often require decision-making based on what has happened earlier in a simulation, suggesting a requirement for memory and logical operators.)

Validation at this level helps to avoid an excessively simplified view, for example that crowds are like seeds in a hopper, because any such position would need to be introduced as an implementation detail rather than as an explanation of behaviour at the reduction level. Once we note at the reduction level that people are not like seeds, the selection of this type of implementation seems forced. Instead, the reduction level's focus on the individual perspective within the crowd suggests modelling of individuals as agents with simple cognitions. This interaction between the implementation and the reduction avoids exclusion of the human element from the simulation; this makes the results more applicable to humans, increasing its validity according to Sime's concerns.

Reduction. The purpose of this level is to develop simple local rules that provide an abstract mechanism for the behaviour as described in the specification. The validation question at this level, then, concerns whether the simple local rules actually implement the behaviour described at the specification level. In other words, are tests run with the local rules essentially tests of the

theory proposed at the specification level? Although there may be multiple ways to reduce the specification to simple local rules, as discussed in section 4.2.2 we are less concerned with the internal mechanism of agents and more concerned with behavioural repertoire; validity at this level, then, concerns whether the behaviours generated by the microscopic human factors are sufficient to implement our theory of how crowds work in the world (the specification).

Our trivial example (figure 3) illustrates another point relating to the validity of the reduction. In many cases the behaviour of the specification will be richer than we wish to model in a particular abstraction, resulting in a gap between a richly detailed specification and an abstract reduction. (In the example it was judgements about cashier efficiency that were seen to be interesting but extraneous to the purpose at hand.) This gap is intrinsic to abstract modelling — as is suggested by section 4.2.2 — and goes to the heart of the general challenges in validating microscopic models described above; we must create abstractions to limit the scope of a phenomenon and achieve understanding, but these abstractions must be defended. The art of abstract modelling is in selecting the right abstraction, a process that must be confirmed as described in the previous paragraph. Many such gaps, however, have a practical benefit: that of identifying areas of interesting future work — a principle that will be well-illustrated in the demonstrations of chapters 6–8. In the limit, these pieces of future work can be modelled, and their effects on the basic behaviour confirmed.

Specification. The purpose of the specification level is to describe a theory or hypothesis of how a particular aspect of crowd behaviour unfolds in the real world. Thus, validation of the model at this level implies a determination of the extent to which the model is a representation of the behaviour of real crowds. Bearing in mind the discussion of validity and microscopic models (above), let us consider how MHF models are relevant to real crowds.

The validation task, to be carried out through experimentation upon the model, is to show that the specification is *sufficient* to explain the real world behaviour of interest. Sufficiency is stressed here, because the specification is subject to the principles of multiple realisability: There may be multiple versions of the specification that can produce the behaviour. To discuss how a model that expresses sufficiency can tell us about the real world, let us consider two cases for

application of the methodology. (Here we assume that, through experimentation on the model, the sufficiency of the specification has been shown.)

In the first case we use the methodology to study a behaviour that is well understood. In this case, the specification is motivated by — or derives directly from — existing knowledge. Here the modeller can make an argument that the specification is more than sufficient to account for the behaviour because the existing knowledge provides a reason to prefer this particular specification over other possible mechanisms. The model can then be used to draw conclusions about how the real world works, supported both by the MHF methodology (that makes the claims about the world explicit) and the literature (or other reasons, that cause us to believe these claims are well-supported).

In the second case, the methodology is used to study a behaviour that is suspected, or to investigate ‘what-if’ scenarios. Here, the MHF model’s contributions are more theoretical in nature. As noted above, microscopic models are generally focused on the qualitative study of emergent effects; they are one tool to help in understanding human behaviour in crowds. Like in much scientific endeavour, particularly in the social sciences, we look for corroborating evidence from multiple disciplines and approaches to answer complex questions; individual methods, such as our methodology, generally contribute by narrowing the field of inquiry rather than through definitive answers. In this respect, the methodology proposed, with its emphasis on local mechanisms, can contribute by demonstrating the sufficiency of individual-based human-centred hypotheses for crowd effects.

In both cases, the model adds to our knowledge of crowd behaviour by testing the sufficiency and consistency of a theory to reproduce crowd effects. This allows us to make in-principle arguments about how crowds can and cannot work. Where we wish further certainty, we can use the narrowed focus of the model to produce hypotheses that can be subsequently tested in real crowds, or — if the hypotheses cannot be practically and/or ethically tested — make philosophical arguments in support of their applicability to real crowds. Additional support for the theory may derive from case studies of natural experiments, and the results of other disciplines.

Although multiple realisations of a particular crowd behaviour may be possible — as we have seen in previous chapters — those that take account of Sime’s concerns are preferred. The advantage of the MHF methodology is that it provides an explicit mechanism for taking account of Sime’s concerns at the heart of a microscopic implementation by providing simple local rules for crowd behaviour. This, then, is a key accomplishment of the MHF methodology: By taking account of Sime’s concerns, we stop modelling ball-bearings and improve the relevance of microscopic models to people. With this methodology, and the simulations derived from it, we are in a position to better understand what emergent behaviours are likely, and to study how human behaviours interact in a large-scale movement simulation.

Exploring the mechanism

Microscopic human factor models have much to contribute that has already been discussed from the point of view of validation. There is, however, more to modelling than the simple reproduction of emergent crowd effects.

These additional benefits are obtained if a *correspondence* is made between abstract rules at the reduction level and the formal structures that implement these rules at the implementation level. In so doing, the model provides a locus for experimentation with the reduced rules by allowing the modeller to modulate or even deactivate some aspects of the simulation. This experimentation can lead to an understanding of how and why the rules provide for emergent crowd level effects as observed in the simulation, and how the new rules interact with previous behaviours of the model. By extension, this provides an opportunity to determine how varying the rules and the situation can give rise to different crowd behaviours, and hence some idea of the range of potential crowd behaviour.

4.4 Summary

Microscopic Human Factors

The microscopic human factors approach aims to integrate human behaviour into microscopic models to satisfy Sime's concerns that behavioural and physical factors are not sufficiently integrated in crowd research. Having rejected the cognitive architecture approach in the microscopic context, we propose a method centred around increasing behavioural repertoire of agents as an interaction between the environment, the individual's social and physical context and the decision making process. We focus on generating behaviours that underlie crowd effects, rather than fidelity to processes internal to the individual.

In modelling a crowd behaviour, the methodology requires three descriptions in increasing levels of abstraction (see table 1). At the specification level, a clear description of how a crowd behaviour occurs in the world is presented, and it is at this level that arguments are made concerning the reasonableness of this starting position. At the reduction level, the specification of how the world works is reduced to a series of abstract individual-centred rules that provide for the behaviour. All modelling and implementation details are relegated to the implementation level, at which the reduced rules are translated into a specific formalism for simulation.

The levels of the methodology support validation concerns at each level. These relate to formalism appropriateness, implementation fidelity, sufficiency of the rules and culminate in an argument concerning the sufficiency of the specification and its role in supporting understanding of crowd behaviour with other disciplines. Models built according to the MHF methodology

Table 1: Summary of MHF methodology levels, their content, purpose and value in validation

| Level | Design Goal | Vocabulary Used | Validation Goal |
|----------------|--------------------------|-----------------------|---|
| Specification | Theory about world | People in crowds | Theory sufficient to account for real-world behaviour |
| Reduction | Essential situated rules | Individuals in crowds | Rules sufficiently capture essence of specification |
| Implementation | Integration with model | Formalism-specific | Implementation corresponds to rules |

(with the additional requirement of correspondence between the reduction and implementation) support experimentation to identify how the modelled rules bring about the behaviour of interest, and variation in crowd behaviour.

The levels also highlight that different objections to a model are possible at different levels. Objections at one level may not invalidate work at another level. For example, there may be legitimate differences in opinion between domain experts on the very questions that must be answered to make an implementable model. This methodology localises this disagreement at the specification level. By creating a different specification for each position, satisfactory reductions of each position to rules is possible, and these models may shed light on the disagreement itself. In other cases, the objections may be more at the level of implementation, with generally agreed specifications and reduction but criticism that modelling simplifications introduced changes to the rules. Isolating these issues to particular levels of the methodology results in increased clarity in the academic debate.

Research objective, relevant behaviour and the goals of the methodology

We began this chapter — inspired by the problem identified in section 3.3 — with our research objective:

Research objective: *To develop a methodology to integrate relevant human behaviours with physical factors at the heart of microscopic crowd models for the purpose of improving relevance to people and studying emergent crowd-level effects.*

We identified a lack of behavioural repertoire as our focus in the development of our methodology. We developed a methodology whose *key goal* is to represent human behaviour as simple local rules at the heart of microscopic models, to support models that can take account of Sime's concerns and shed light on the mechanisms of emergent crowd effects.

We shall present concrete examples of the methodology at work in chapter 6–8, and we shall discuss the performance of the methodology in these demonstrations, as part of the general discussion, in section 9.2. In the demonstrations, a consistent underlying formalism will be

used: that of the floor field model. Accordingly, before we turn to the demonstrations of the methodology, we present the details of the floor field model's implementation.

Chapter 5

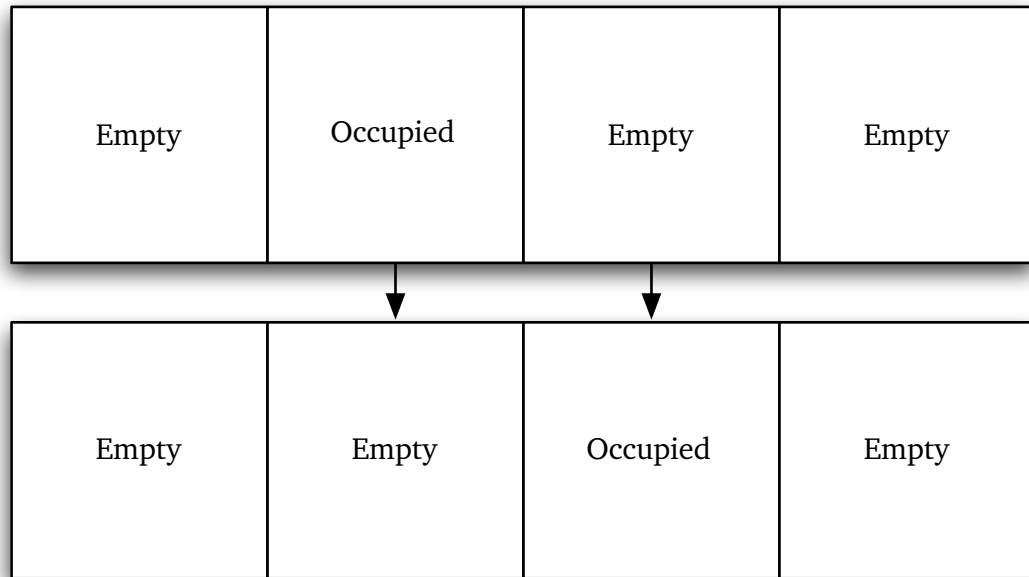
The floor field model

As previously discussed, we have found the floor field model cellular automaton to be an excellent formalism for implementing models that use the Microscopic Human Factors methodology. We will be using this formalism in chapters 6–8 for all three of our demonstrations of the methodology. In order to clearly express work based on the floor field model, a description of the base model is required; the purpose of this chapter is to provide such a description.

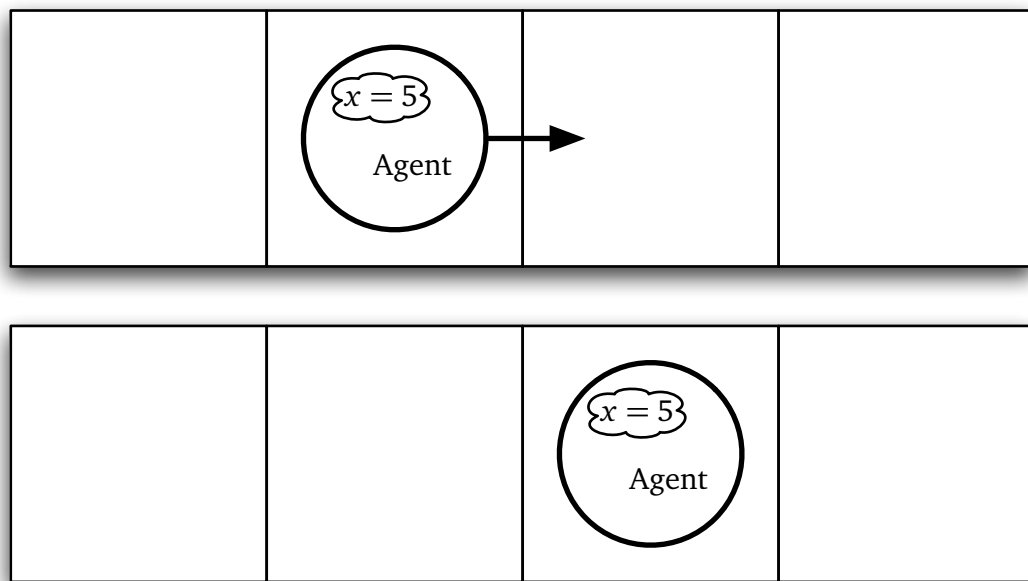
The floor field model (Burstedde et al., 2001) has proven to be full of potential, and has been the subject of a number of refinements and extensions. As such it is necessary to select a point in time that serves as a baseline for our work. We have adopted the model as it was presented in an exploration of the floor field model’s applicability to egress simulation (Kirchner & Schadschneider, 2002), which allows us to relate our egress results to a comparable characterisation of the floor field model.

5.1 Cellular automata and Multi-agent systems

Formally speaking, the ontology of cellular automata does not include mobile agents; as in figure 4a, grid cells only exist in a particular state (here, either occupied or empty). Clear rules operate on a cell-by-cell basis to specify when a cell will change state. What is sometimes termed *movement* in a cellular automaton is more formally described as a transition between states in two adjacent cells (here, from occupied to empty in one cell, and from empty to occupied in



a. Movement in a cellular automaton. Two separate state transitions: Left cell gets Occupied \rightarrow Empty and right cell gets Empty \rightarrow Occupied.



b. Movement in a multi-agent system. Note mental content (here, $x = 5$) moves with agent.

Figure 4: Comparison of state and movement in a cellular automaton and a multi-agent system. In this case, space is one-dimensional. Upper boxes in each sub-figure represent time $t = 0$, lower boxes represent $t = 1$.

the adjacent one). As information is not copied between cells in a cellular automaton, formally representing complex internal state in a “moving” agent is non-trivial.

By contrast, a multi-agent system of the type described in our work retains the discrete geometry of the cellular automaton and introduces an additional concept: agents that have encapsulated internal state (analogous to mental content). On this view, the grid provides co-ordinates tracking agent position, and allows the assertion of exclusive occupancy of a cell. Location, then, is seen as a property of the agent, and movement naturally preserves internal state (figure 4b).

Despite these formal differences, it is possible to conceive of the floor field model as a multi-agent system rather than as a cellular automaton, as the concepts of the model are well-suited to such an expression. As discussed in section 1.3.2, the individual focus of the multi-agent approach is a natural fit for modelling social systems, and we have found this approach to be a good match for the Microscopic Human Factors methodology. Thus, we here describe the floor field model in multi-agent terms, rather than as a cellular automaton.

5.2 Design of model

The floor field model, according to our agent-based restatement, is a microscopic crowd model of individuals. Agents are initially distributed at random on a 2-dimensional grid that provides a co-ordinate system, both for movement and for maps of information available to agents called fields. Agents move according to local rules, balancing their movement decisions between reducing distance from desired goals and following other nearby agents. The floor field model as originally specified does not provide for individual cognition beyond combining the two perceptions just described. (A complete algorithm for the model will be presented on page 77.)

5.2.1 Physical environment

Movement. In the floor field model, agents have a single action: they move. The discrete time evolution of the model thus considers the dynamic and unfolding occupancy of cells. Agents move a distance of zero or one cell per time step, and movement occurs in parallel. Because cells

can hold at most one agent, this parallel update ensures that if two agents attempt to move to an unoccupied cell a random agent will succeed. (For an alternate mutual exclusion procedure, see Nishinari et al., 2004.) Due to the cell selection rules (see below), in the floor field model agents only try to move to unoccupied cells.

Walls and exits. Certain cells within the space (e.g. around the perimeter of the modelled area) can be designated as wall cells. These cells are unavailable for movement. For the purpose of modelling ingress or egress behaviours, the floor field model contains designated exit cells, generally within the walls. Agents move onto these cells as usual and occupy them for the duration of the time step. The next move of an agent occupying an exit cell is to disappear from the model.

Cell adjacency. The model uses the Cartesian directions in determining cell adjacency. In other words, allowable movement directions on the two-dimensional grid are North, South, East and West. (In the floor field model it is often required to speak of the possible cells an agent could move to in the next time step. This is called the *neighbourhood* and includes the four adjacent cells, as well as the agent's current cell.)

5.2.2 Floor fields

The floor field model represents information available locally by distributing it throughout the space. This can take the form of continuous information (similar to a temperature measurement at each cell) or discrete information (similar to elementary particles like bosons distributed on cells). By analogy to physics, we can speak of a *field*: a set with one integer or real value per cell (as cells are indivisible). By consulting the field values in the neighbourhood, agents extract localised information about their surroundings. The name *floor field model* derives from these fields, which serve as spatial data structures. Aside from the presence of wall cells or other agents in the neighbourhood, agents do not have recourse to any information that is not encoded in these structures.

Taking temperature as a trivial example of the use of fields, a room may have a gradient of temperatures, detectable as we move through a space. We can find a heat source in the room without knowing its precise position by choosing steps through the space in the direction of increasing heat. This corresponds to an agent considering local values of a floor field, moving in the direction of increasing value.

Static field. The first of the two fields provided by the model is the static field, which represents the agents' mental map of the environment (see figure 5). The static field — in an in/egress scenario — is a continuous field that encodes the distance from a cell to the nearest entrance/exit. By using the information in the static field, agents can determine whether a step in a certain direction will bring them toward an exit. During egress, the static field can be considered a measure of desirability; cells closer to an exit are more desirable. In this sense, the local value of the static field can be thought of as the value of a heuristic function.

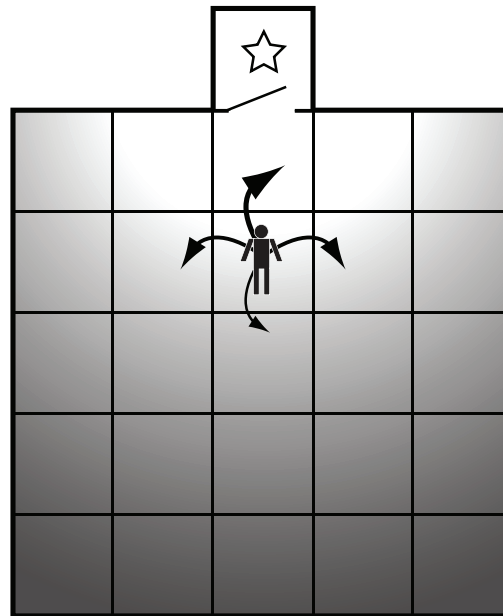


Figure 5: Floor field model space with static field, S . Agent (*pictured*) can move to neighbouring cells (*arrows*) or current cell (*arrow not shown*). Cell selection is probabilistic based on agent's perceived desirability of cell (*arrow size*). Movement to exit cell (*starred*) results in subsequent disappearance of agent. Static field (*gradient overlay*) represents inherent cell desirability; field strength on own and neighbouring cells is accessible to the agent.

As expected from the name, the static field is fixed from the start of the simulation run, and does not change during the run. As is common with heuristic function definitions, the field's concept (distance from exit) can be quantified according to one of several metrics, with different simulation dynamics resulting (Nishinari et al., 2004; Kretz, Boenisch, & Vortisch, in press). The present incarnation of the floor field model uses a simple linear distance measure for the static field (Kirchner & Schadschneider, 2002), which we define as follows:

$$s(i, j) = \min_{(d_x, d_y) \in X} \sqrt{(d_x - i)^2 + (d_y - j)^2} \quad (2)$$

$$s_{max} = \max_{\forall (i, j)} s(i, j) \quad (3)$$

$$S_{ij} = s_{max} - s(i, j) \quad (4)$$

Here $s(i, j)$ gives the distance from a cell with co-ordinates (i, j) to the closest exit in X , the set of co-ordinates of all exits. The normalisation number, s_{max} , is the greatest distance to an exit from any cell. The final static field value, S_{ij} decreases from s_{max} at exit cells to 0 on the furthest cell from an exit.

Dynamic field. The second of the two fields provided by the floor field model is the *dynamic field*. It provides a mechanism for agents to become aware of the movement of others by analogy with ant pheromone chemotaxis (c.f. Bonabeau et al., 1991). It is a discrete field, whose integer value on each cell is the number of bosons located on that cell (initially zero). When an agent moves from a cell (i, j) to an adjacent cell it drops a boson on the origin cell; this is represented in the dynamic field by $D_{ij} \rightarrow D_{ij} + 1$. Thus, the boson acts as a sort of breadcrumb trail that can be followed by other agents (see figure 6).

The dynamic field is so-named because bosons have dynamics by which they diffuse to neighbouring cells and decay. During each time step, each boson in the model decays (vanishes) with probability δ . Those particles that do not decay may diffuse (move to a randomly selected adjacent cell) with probability α . Both of these probabilities are parameters to the model. Agents can consult the values of the dynamic field in neighbour cells in order to follow virtual 'paths' left by previous agents. Due to the effect of δ , paths must be constantly refreshed in order to be

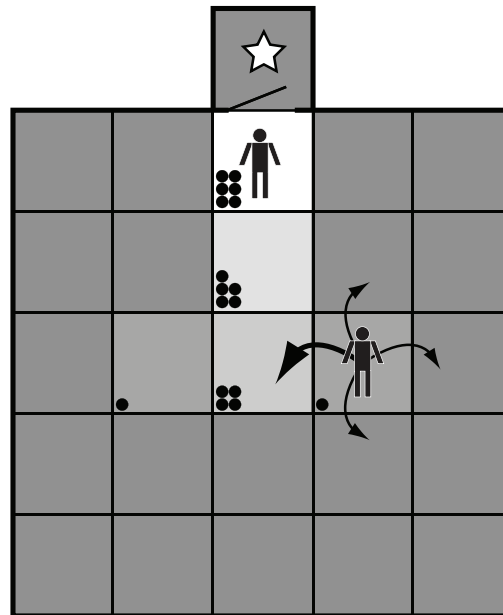


Figure 6: Floor field model space with dynamic field, D . Previous agents (e.g. *top agent moving up from centre*) have dropped dynamic bosons (*circles*). Bosons probabilistically decay (*middle group*) or diffuse to neighbouring cells (*lower group*). An agent attending only to the dynamic field uses boson density like a trail (*lower agent following gradient overlay*) to follow preceding agents.

effective at guiding pedestrians, while α has the effect of spreading information about a popular route to surrounding cells.

While ant pheromone chemotaxis is a measurable phenomenon, the floor field model does not suppose that humans are guided by chemical secretions in crowds. Rather, the dynamic field is an analogy adopted in order to make the simulation more efficient by avoiding long-range interactions between agents (Burstedde et al., 2001). By using the diffusion and decay of the dynamic field, long-range interactions between the agents can be efficiently reduced to local ones. Just as the decay and diffusion of ant pheromone odours allow for ant trails become optimised and furthermore to change dynamically as required, the dynamic field allows agents to co-operate in path finding in the model.

5.2.3 Cell selection

To guide agents through the model, a score is assigned to each cell indicating its desirability:

$$c_{ij} = \exp(k_D D_{ij}) \exp(k_S S_{ij}) (1 - n_{ij}) \xi_{ij} \quad (5)$$

Here, n_{ij} is 1 for occupied cells (otherwise 0) and ξ_{ij} is 0 for walls (otherwise 1); this ensures that walls and occupied cells will not be selected by agents for movement.

Two sensitivity parameters, k_S and k_D are provided as parameters to the model, as it is desirable to regulate the information available to the agent from each of the floor fields. (It may be that the agents do not have perfect information concerning the movement of others, for example due to darkness.) The k_S and k_D parameters scale the influence of S_{ij} and D_{ij} (the static and dynamic field values, respectively); a sensitivity parameter can increase a field's influence ($k > 1$) or decrease it ($0 \leq k < 1$). To take an egress example, if k_S is low, then the agent moves through the grid ignorant of where the exits are. If k_S is high, then the agent is attuned to the location of exits. If k_D is low, then the agent is not concerned/aware of the movements of others. If k_D is high the agent will be disposed to follow other agents through the grid.

Agents engaged in selection of a target cell for movement convert the scores on neighbouring cells to probabilities, according to the following equation:

$$p_{ij} = \frac{c_{ij}}{\sum_{(i',j') \in N} c_{i'j'}} \quad (6)$$

Here, N is the set of co-ordinates of neighbour cells. Each agent probabilistically selects a cell in this manner in every time step. All agents select cells at the same time and before any agent attempts movement. Subsequent movement conflicts are resolved according to the details provided above.

5.3 Sensitivity parameter effects

A moment of reflection is in order upon the effects that k_S and k_D have on the set of probabilities p_{ij} in a neighbourhood. These probabilities always sum to 1, and the normalisation process

of equation 6 creates the probabilities with the *available* information provided by equation 5. This normalisation, and hence agent preference, is guided by trade-offs between k_S and k_D in ways that may be initially non-evident. Consider an isolated agent in the centre of a room. If the agent is prevented from accessing the *static* field ($k_S = 0$) then the value of p_{ij} will be solely influenced by the *dynamic* field. If, however, the static field's influence is increased to have equal importance ($k_S = k_D$) then normalisation must decrease the importance of the dynamic field to balance the effects of the fields. The point here is that the normalisation of scores to probabilities means that lowering one sensitivity parameter has the effect of increasing the influence of the other.

It is not the case, however, that probabilities in the system are due only to the ratio $k_S : k_D$. The exponentiation of the k -adjusted field values allows variation of the sensitivity parameters to emphasise or de-emphasise small differences in field values between neighbours. If $k_S = k_D$ but both are high, for example, then small differences in either field between neighbours become large differences in the cell score and larger probability spreads are obtained. At the other extreme, if $k_S = k_D = 0$ then differences between fields do not matter at all as $c_{ij} = 1$ for all empty non-wall cells and the probability of moving to any of the five neighbours will be 20%. We shall return to these points in chapter 7.

5.4 Model evolution

The floor field model simulation begins with the creation of the static field according to equations 2–4, and the distribution of agents through the space. Thereafter, the model proceeds in a loop repeating the instructions for the basic time step, given in procedure `STEP` in algorithm 1. In this algorithm, the agent methods `SET-DESIRED-CELL`, `GET-DESIRED-CELL` and `GET-CURRENT-CELL` are simple accessors whose implementation is not shown. The trivial procedures `DECAY` (cause a dynamic boson to evaporate), `DIFFUSE` (cause a dynamic boson to move to a random adjacent cell), `SHUFFLE` (randomly re-order a list), `MOVE` (cause an agent to change position on the grid), `GET-NEIGHBOURS` (return a list of adjacent cells) and `RANDOM` (return a random number in the range $[0, 1]$) are considered primitive and their implementations are not shown. Although no

use is made in algorithm 1 of the return value from `ATTEMPT-MOVE`, future algorithms will make use of it so it is shown here for consistency.

5.5 Summary

We have presented the details of the floor field model (Burstedde et al., 2001, as modified in Kirchner & Schadschneider, 2002). The model has been restated as a multi-agent system in order to accommodate state (i.e. memory, experience) the modeller may choose to define; any such state will be carried along with the agent in its movements through the model.

We shall rely heavily on the contents of this chapter in explaining the implementation and results of our demonstrations. Algorithm 1 will be elaborated in each demonstration to show how we integrate the human behaviour simulations tightly with the basic model as described in this chapter.

We turn now to the first demonstration of the methodology, in which we model a simple agent behaviour (voluntary pushing) as well as a simulation of force and crowd safety, in the floor field model.

```

STEP( $D, S, agents$ )
1  DECAY-AND-DIFFUSE( $D$ )
2  for each  $a$  in  $agents$ 
3      do  $a.CELL-SELECT(D, S)$ 
4  for each  $a$  in SHUFFLE( $agents$ )           ▷ SHUFFLE simulates parallel update
5      do  $a.ATTEMPT-MOVE()$ 

DECAY-AND-DIFFUSE( $D$ )
1  for each  $boson_{ij}$  in  $D$ 
2      do if RANDOM() <  $\delta$ 
3          then DECAY( $boson_{ij}$ )
4          elseif RANDOM() <  $\alpha$ 
5          then DIFFUSE( $boson_{ij}$ )

AGENT METHOD CELL-SELECT( $D, S$ )
1   $s \leftarrow SELF$ 
2   $N \leftarrow s.GET-NEIGHBOURS()$ 
3   $r \leftarrow RANDOM()$ 
4   $w \leftarrow 0$ 
5  for each  $(i, j)$  in  $N$ 
6      do  $w \leftarrow w + p_{ij}$ 
7      if  $w < r$ 
8          then  $s.SET-DESIRED-CELL(i, j)$ 
9      return

AGENT METHOD ATTEMPT-MOVE()
1   $s \leftarrow SELF$ 
2   $(i_0, j_0) \leftarrow s.GET-CURRENT-CELL()$ 
3   $(i_1, j_1) \leftarrow s.GET-DESIRED-CELL()$ 
4  if IS-EMPTY( $i_1, j_1$ )
5      then DROP-BOSON( $i_0, j_0$ )
6           $s.MOVE(i_1, j_1)$ 
7      return TRUE
8  else return FALSE

```

Algorithm 1: Floor field model time step function

Chapter 6

Voluntary pushing, force and crowd safety

We now turn to the demonstrations of the methodology in use. The present chapter and the two following chapters implement three broad crowd behaviours. By using the methodology to study these crowd behaviours we hope to accomplish three things: first, to show that the methodology is practical; second, to show that using the methodology we can build models that take account of Sime's concerns; and third, to show that models built with the methodology have value and can generate interesting discussion and conclusions. Each chapter will highlight an application of the methodology to a different type of situation, building upon the previous chapters, and using the floor field model as an implementation formalism. This chapter focuses on a known crowd behaviour, where our goal is to add a simple individual behaviour to agents and investigate and characterise the emergent crowd-level effects that ensue. Later chapters will consider changes intended to directly affect the crowd level itself, and to study a more prospective hypothesis.

The first demonstration concerns a simulation of the application of physical force in crowds. The demonstration of force in this chapter differs from the remaining chapters in that there is more to integrate with the floor field model than just human behaviour. The floor field model is lacking a simulation of force — crucial to consideration of crowd safety, and also to its ability to represent the pushing behaviours that we want to add. Although the Microscopic Human Factors methodology was conceived as a method for adding human behaviour to microscopic models, it is equally well suited to microscopic implementation of individual-based physical factors such

as force. Accordingly, in this chapter we integrate both human sources of force and the force simulation itself into the floor field model using the Microscopic Human Factors methodology.

As discussed in chapters 1 and 2, force plays a major role in crowd disasters. It is the hypothesis of this chapter that physical force is characteristic of crowds, and cannot be abstracted out of the modelling of dense crowds without a dramatic alteration in crowd dynamics. In the case of a crowd exiting through a doorway, we expect that significant changes to numbers of agents exiting, exit time and exit dynamics would be observed. In addition, without a concept of force, a model cannot predict the injuries that directly affect the injured, and indirectly affect other individuals trying to exit around them.

If forces are characteristic of crowds and crowd disasters, we might ask what level of commitment a model ought to have with respect to forces. We see three possible positions for treatment of force in a model:

1. The first position (taken implicitly by the floor field model) hypothesises that force is not an essential characteristic of a crowd model; the corollary to this position is that results can be accurate with or without force so there is little lost by abstracting it out of a model.
2. The second position is that force can be accounted for without an explicit simulation (perhaps as a linear scaling factor upon the output of a non-force model).
3. The third position, which we here espouse, is that, being characteristic of crowds, force has a complex effect on individual actions; force cannot be removed from a crowd model without totally changing the dynamics of the system.

Our plan is to implement force within the floor field model, retaining its individual focus, particle field metaphor and the spirit of its movement rules. Through an experimental study of its effects, we shall determine which of these positions is supported by the evidence. We shall return to this question at the end of the chapter.

The three level methodology requires three distinct descriptions of physical force application: *specification*, *reduction* and *implementation*. Each of these will be treated in sequence before we turn to results and discussion.

6.1 Specification: Forces in crowds

In section 2.1.2 we gave an abridged citation of Fruin's characterisation of forces in crowds. A more complete excerpt will help to explain our conception of crowd forces:

It is difficult to describe the psychological and physiological pressures within crowds at maximum density. When crowd density equals the plan area of the human body, individual control is lost, as one becomes an involuntary part of the mass. At occupancies of about 7 persons per square meter the crowd becomes almost a fluid mass. Shock waves can be propagated through the mass sufficient to lift people off of their feet and propel them distances of 3 m (10 ft) or more. People may be literally lifted out of their shoes, and have clothing torn off. Crowd forces can reach levels that are almost impossible to resist or control. Virtually all crowd deaths are due to compressive asphyxia and not the "trampling" reported by the news media. Evidence of bent steel railings after several fatal crowd incidents show that forces of more than 4500 N (1000 lbs.) occurred. Forces are due to pushing, and the domino effect of people leaning against each other (1993).

Several points of interest can be gleaned from this description of extreme crowd forces. First, Fruin implies that people have individual control only under certain circumstances, namely that of low density and low force application. Second, Fruin highlights the involuntary nature of movements in dense forceful crowds; this recalls Johnson's description (see p. 22) of the inability of those near the pile of people at Riverside Coliseum to avoid standing or falling upon others. Third, force propagates throughout a crowd (like a shock wave). Fourth, crowd deaths are due to force. Fifth, forces originate through individuals pushing, and are compounded through leaning forces (we call the former source *voluntary pushing* and the second *involuntary pushing*).

If damaging crowd forces originate with voluntary pushing, we may well ask why do people in crowds push in the first place? We are not aware of a formal study of this basic question, but the individual focus of the methodology requires that we take a position on the topic if we wish to model it as an individual behaviour. Although Johnson reported on young concert-goers experimenting with shock wave creation, we suspect this is not the primary motivator for voluntary application of force in crowds. The Bethnal Green disaster (see section 2.2.1) is perhaps a more clear example where force was applied where people were frustrated in the slow pace of movement to a safe area. This resonates also with crazes in which people are competing for a limited resource, and pushing could occur due to frustration in achieving goals. Because

of the uncertainty concerning reasons for force application in crowds, in our analysis we focus more on the effects of force once present, rather than making claims about force origination; we take the position that frustration of movement goals results in force application.

Fruin's citation certainly highlights the important role that force plays in crowd disasters. It is not only the obvious lack of injuries that results from the absence of force in a model. The loss of individual control means that as people move within a crowd, force effects can cause them to be diverted from their individually planned movement pattern. We suspect that this diversion causes a change in exit time that cannot be accounted for in the absence of force. In addition, injuries, a direct consequence of force, themselves have effects on non-injured people still trying to exit; will such effects be obstructive or facilitative? In either case, injuries may affect exit times in a way that would not be possible for a force-free model to explain. By exploring the specification we have presented, we hope to shed light upon these questions.

6.2 Reduction: Local rules for force behaviours

The purpose of a model is to reduce the complexity in a real-world system down to its essential aspects. The role of the reduction is to identify the rules that we feel are essential from the specification. Accordingly, in order to incorporate force and pushing behaviours into our crowd model, we determine those essential aspects of force that must be represented. The language of the reduction level avoids formalism-specific terminology or modelling simplifications, and adopts the local focus that is the perspective of individuals within the crowd. As discussed at the start of this chapter, we take a similar approach to developing the physical force simulation itself. We have adopted four essential principles concerning the modelling of voluntary pushing and force in crowds:

1. Force is *generated purposefully*. Force is not simply an emergent property of masses of people. People exert force in some circumstances and not in others. In the specification we proposed that frustration initiates force; at the level of the reduction we look for an abstract mechanism to stand in for the concept of frustration in triggering the application

of force. We adopt the rule that agents push when a desired movement is blocked by the presence of another agent.

2. Force is *directed*. It is applied by one person to another in a particular direction.
3. Force is *propagated*. Fruin tells us that forces in crowds are due to a “domino effect” of pushing and leaning. This means that force in a crowd model must be transmitted through the crowd in the direction of the force, and that it is additive (vector-wise) when encountering other forces. Moreover, Fruin’s reports of shock waves imply that propagation occurs over a certain time period. Force applied at the rear of a crowd is not immediately felt at the front; instead it travels from person to person, and is experienced by intermediate people as it is transmitted.
4. Force carries *consequences*. Although force can have an effect by guiding cognitive decision making (see chapter 8 for more on this point), in this chapter we are concerned with its physical consequences. Once exerted, force has measurable effects that are outside a person’s control. Fruin suggests that people can be overcome by excessive force, but also that force can result in a loss of individual control. Thus, from an individual perspective, a person who is being pushed is caused to move and/or lean in the direction of the force, not in any other direction they may have desired.

People who experience excessive force become injured and no longer move within the model. Although it seems obvious that a range of injuries can be produced in real crowds (including death by compressive asphyxia at the extreme), for our abstract purposes we define injured agents to be immobilised people who simply take up space. We shall return to this point in section 6.9.

6.3 Implementation

We have implemented the rules described by the reduction at the heart of the floor field model (Burstedde et al., 2001). We have sought, wherever possible, to be faithful to the design patterns inherent in the floor field model, in order to allow comparisons with the basic model. As

described in chapter 5, we have adopted the version of the model from (Kirchner & Schadschneider, 2002) for the purposes of comparison with the data described in that paper. Our changes can be divided into three categories: addition of a dynamic force field, cell selection changes and addition of injuries. These changes will be fully described first in textual form, but are summarised in pseudo-code in algorithm 2 on page 87.

6.3.1 Force Field Addition

The underlying change that supports simulation of individual pushing in the floor field model is the addition of a dynamic force field. In keeping with the floor field model's physical perspective (agent interaction with elementary particles through static and dynamic fields) forces are represented in the model through force bosons. Force bosons are unit vectors, analogous to directional elementary particles. The force field has a discrete value on each cell, namely the aggregation (vector sum) of the force bosons on that cell. Thus the force field represents the direction and magnitude of the force experienced by an agent on a cell. The force on a cell is not perceptible by agents on neighbouring cells.

Force is generated (force bosons deposited) within the model by agent pushing. Agents push when they desire to move to a cell but cannot do so because it is occupied. Agents within the model have a variable capacity for pushing (ρ , the number of bosons dropped per push) that is selected at the start of the model from a normal distribution whose mean and standard deviation are input parameters to the model.

Each time step, the force bosons on a cell are propagated to neighbouring cells according to Fruin's observation (discussed above) that crowd forces are repeatedly re-transmitted from person to person through interpersonal contacts within crowds. The consequence of this re-transmission at the agent level is that the individual force bosons are not necessarily preserved. (In a crowd an individual pushed with equal force from the left and right takes up both forces; the rightward force is not retransmitted through the individual to the other side.) Agent-level retransmission means that it is the aggregated force field vector (\vec{f}_{ij} of equation 7, in which B_{ij} is the set of force bosons on cell i, j) that is propagated, rather than the individual bosons

themselves. Force is dissipated (i.e. its underlying bosons disappear) if it would move onto an empty cell (one in which $n_{ij} = 1$ in equation 5) or a wall cell (one in which $\xi_{ij} = 0$ in equation 5), or a cell with an injured agent (see below).

$$\vec{f}_{ij} = \sum_{\vec{b} \in B_{ij}} \vec{b} \quad (7)$$

The propagation of \vec{f}_{ij} occurs by depositing $|\vec{f}_{ij}|$ force bosons onto the cell in the direction \hat{f}_{ij} . Because the coordinate system is discrete while \vec{f}_{ij} is continuous, both the magnitude and direction of \vec{f}_{ij} must be quantised. Accordingly, $\text{floor}(|\vec{f}_{ij}|)$ is used for the magnitude, while a probabilistic method is used to quantise \hat{f}_{ij} to one of the four neighbouring cells. When \hat{f}_{ij} points between two cells, the model — in considering each force boson to be deposited — assigns the boson to one of the two cells probabilistically. The probability of being deposited on one of the cells is inversely proportional to the deviation of \hat{f}_{ij} from that cell and is given by equation 8, in which p_a is the probability of selecting the neighbour with the lower angle and p_b is the probability of selecting the neighbour with the higher angle.

$$1 - p_a = \frac{\hat{f}_{ij} \bmod 90}{90} = p_b \quad (8)$$

6.3.2 Changes to agent behaviour

Impact of force. While the force simulation calculates the force on a cell using a vector sum, the agent's experience of force is calculated using a scalar sum for the purposes of cell selection and injuries (see below). This means that two equal and opposing forces that cancel each other out on a particular cell will still have an impact on an agent on that cell who becomes squeezed between those forces. Where required we shall use the symbol \hat{f}_{ij} rather than \vec{f}_{ij} to denote this scalar sum of forces.

Cell Selection. The floor field model's cell selection mechanism has been changed in two ways. Most importantly, the model is changed so that normal cell selection is bypassed if the force experienced by an agent exceeds a threshold, χ (we have used 3ρ , three times the force the agent can produce when pushing others). If this level of force is experienced, the agent is

deemed to have selected the cell in the direction of the force. Quantisation using equation 8 is performed as required. With two exceptions, other model rules and behaviour have not been changed. In particular, the agent is not permitted to share a cell with another agent or a wall, movement is still simultaneous, the agent will push if the move cannot be made (the latter represents the involuntary, or leaning forces within the model). The two exceptions are required in order to make the force rules effective:

First, propagation depends on a crowd in which agents vie for empty space and in which movement in and out of cells can be simultaneous. This is because force is not propagated through empty space. In the floor field model, however, equation 5 requires that cells artificially lie empty for at least one iteration after being occupied. This occurs because the term $(1 - n_{ij})$ of equation 5 drops the probability of selecting an occupied cell to 0, preventing agents from selecting an occupied cell to move into it as the previous agent leaves it. This artificially spaces the crowd out, preventing effective force propagation. We are unaware of any particular basis for this property of the floor field model, and so we propose a refinement to the behaviour. Although we could have simply removed the $(1 - n_{ij})$ term, we feel it is reasonable that agents will look favourably upon empty cells when others are occupied (the choice of an occupied cell in effect being a bet that the occupying agent will move on). Accordingly we replace the occupancy number n_{ij} of equation 5 with our vacancy factor η_{ij} which is 0 if cell (i, j) is unoccupied and 0.5 otherwise. This term halves the cell score for occupied cells relative to unoccupied ones.¹ The revised cell selection formula appears as equation 9.

$$c_{ij} = \exp(k_D D_{ij}) \exp(k_S S_{ij}) (1 - \eta_{ij}) \xi_{ij} \quad (9)$$

Second, the introduction of the η_{ij} parameter to remediate the problem just described is incomplete without a switch to the use of the more familiar (for multi-agent modelling) 4-cell

¹Note that it is the cell score that is halved, as opposed to the probability of selecting that cell (which is derived from the cell scores by normalisation). The consequent effect on probability for an occupied cell is dependent on the distribution of cell scores in the neighbourhood. Consider the example of an agent surrounded on three sides by walls, and on the fourth by an occupied floor cell: The wall cells each have a cell score of 0 (due to the ξ term); halving the cell score of the only remaining choice does not affect the agent's 100% chance of selecting it. With highly asymmetric cell scores (e.g. due to high k values or large field differences) the effect of halving a cell score may have only a limited effect on the probability distribution, a point we shall return to in section 6.5. This is justified, as high desirability of a cell implies high agent willingness to wait to occupy it.

Cartesian neighbourhood. This is because the η_{ij} parameter interacts poorly with the 5-cell floor field model neighbourhood (that includes an agent's current cell as a target for movement): The option to select a self-occupied cell, which is always an acceptable 'move' for the agent (implying $\eta_{ij} = 0$ for the current cell) results in large p_{ij} for the self-occupied cell relative to occupied neighbour cells, and consequently a low probability of selecting the occupied neighbours. The purpose of introducing η_{ij} is therefore frustrated, with agents continuing to eschew moves to occupied neighbours. The authors of the floor field model do not provide a basis for adopting the 5-cell neighbourhood and we have not found a compelling argument in its favour. For these reasons, we adopt the more common 4-cell Cartesian neighbourhood, only considering the four adjacent cells and not the agent's current cell. This allows the η_{ij} term to carry out its function of allowing agents to bet on occupied cells.

The effect of these two changes to the basic floor field model will be considered in isolation in section 6.4.2.

Injuries. An agent becomes injured upon experiencing a force greater than a threshold, ϕ , which is an input parameter to the model. Injured agents become totally inactive, no longer moving through the model, and are treated like wall cells by other agents and the force propagation algorithm.

6.3.3 Time step Progression

In addition to the steps required by the original floor field model, the revised model must propagate force, determine injuries and bypass cell selection. Algorithm 2 shows the revised STEP procedure. In this algorithm all other previously defined procedures remain unchanged.

Several procedures are considered primitive and their implementation is not shown: The NOTE-INJURY procedure removes the current agent from the set of agents and notes the injury at the agent's current location. The PUSH procedure drops force bosons on the agent's desired cell.²

²Actually, the force is dropped on the agent's current cell, allowing the normal force propagation routine to propagate the force to the desired agent's cell at line 8 of procedure STEP. This has the net effect of the action described.

```

STEP( $D, S, F, agents$ )
1  DECAY-AND-DIFFUSE( $D$ )
2  for each  $a$  in  $agents$ 
3      do  $a$ .CHECK-INJURY( $F$ )
4           $a$ .CELL-SELECT-WITH-FORCE( $D, S, F$ )
5  for each  $a$  in SHUFFLE( $agents$ )
6      do if  $a$ .ATTEMPT-MOVE() = FALSE
7          then  $a$ .PUSH
8  PROPAGATE( $F$ )

AGENT METHOD CHECK-INJURY( $F$ )
1   $s \leftarrow SELF$ 
2   $(i, j) \leftarrow s$ .GET-CURRENT-CELL
3  if  $\hat{f}_{ij} > \phi$ 
4      then NOTE-INJURY( $s$ )

AGENT METHOD CELL-SELECT-WITH-FORCE( $D, S, F$ )
1   $s \leftarrow SELF$ 
2   $(i, j) \leftarrow s$ .GET-CURRENT-CELL
3  if  $\hat{f}_{ij} > \chi$ 
4      then  $(i, j) \leftarrow CELL-IN-DIRECTION(\hat{f}_{ij}, i, j)$ 
5           $s$ .SET-DESIRED-CELL( $i, j$ )
6      else  $s$ .CELL-SELECT( $D, S$ )

PROPAGATE( $F$ )
1  for each  $(i, j)$  in  $F$  having a force boson
2      do if IS-EMPTY( $i, j$ ) or IS-INJURY-AT( $i, j$ )
3          then  $\triangleright$  force is not propagated
4          else PROPAGATE-FORCE( $\vec{f}_{ij}, F'$ )
5   $F \leftarrow F'$ 

PROPAGATE-FORCE( $\vec{f}_{ij}, F'$ )
1  for  $i \leftarrow 1$  to floor( $|\vec{f}_{ij}|$ )
2      do  $(i', j') \leftarrow CELL-IN-DIRECTION(\hat{f}_{ij}, i, j)$ 
3          if not IS-EMPTY( $i', j'$ )
4              then DROP-FORCE-BOSON( $F', i', j', i, j$ )

```

Algorithm 2: Step function augmented by pushing and force additions

The `CELL-IN-DIRECTION` function accepts \hat{f}_{ij} and a starting cell's co-ordinates, returning the co-ordinates of the adjacent neighbour cell in that direction. The `IS-INJURY-AT` function determines whether a cell is occupied by an injured agent. The `DROP-FORCE-BOSON` function drops a single force boson on cell (i', j') oriented from cell (i, j) . For simplicity of presentation, we assume that if a subscripted value is passed as a parameter to a procedure, the subscripts are available to that procedure as well.

6.4 Results: Adapting the model for force

In sections 6.4–6.7 we compare the effect of our changes to Kirchner and Schadschneider's analysis of the basic floor field model. (We shall refer to Kirchner and Schadschneider's analysis as the *FFM analysis* and to the analysis of our changes as the *force analysis*.) The purpose of this investigation is to investigate the effects produced as we add force to the simulation, and to demonstrate the results than can be expected from a model built using our methodology. Accordingly we have been guided by the FFM analysis and present results in comparative form.

In this section we present the parameters used throughout the force analysis. We then discuss our replication of the original FFM analysis, as well as considering the effects of the non-force-related changes we have described. Following sections turn to the effects of introducing force and injuries into the model.

The discussion of the force-enabled model is divided into three parts due to a need to change the measure used by the FFM analysis. As simulations in the FFM analysis were run until all agents have successfully exited the modelled space, the measure in the FFM analysis is the number of iterations required for this complete exit. Waiting for a complete exit is ill-suited to a model with injuries because an agent can become injured in front of the door in such a way as to prevent other agents from exiting; if agents are prevented from exiting then the number of iterations required for a complete exit is undefined. We have changed to a measure that counts the number of successful exits in a fixed number of iterations.

We bridge between the FFM analysis measure (we call it the *complete exit* measure) and our measure (the *fixed time* measure) as follows: In section 6.5 we run our model using the complete

exit measure with injuries disabled for validation and comparative purposes; this ensures that a complete exit is possible and allows discussion of the basic effects of force. In section 6.6 we show the same scenario (injuries disabled) using the fixed time measure to provide a reference between measures. Finally, in section 6.7 we enable injuries in our model and use our fixed time measure to test the importance of incorporating injuries directly within the model.

6.4.1 Model Parameters

We have adopted the fixed parameter values that FFM used wherever possible. Except as noted in the text, parameters were set as summarised in table 2. The model was implemented using the Repast 3.0 framework for Java (North, Collier, & Vos, 2006). Model execution (in this and subsequent demonstrations) was on an Apple Mac Pro (two 2.66 GHz dual-core Intel Xeon ‘Woodcrest’ processors, 2 GB system memory) running Mac OS X version 10.4.11, and equipped with the system-standard Java revision 1.5.0_13.

Unless otherwise indicated, fifty replications were run for each parameter combination; graphical results show the mean of these replications for the dependent variable in question.

Table 2: Parameter table with default settings

| Parameter | Value | Description |
|-----------------------------|----------------|--|
| Grid size | 63×63 | Outer boundary of modelled space (includes wall cells) |
| Agent Pop. | 1116 | Starting agent density (30% of available cells) |
| α | 0.3 | Probability of dynamic boson diffusion |
| δ | 0.3 | Probability of dynamic boson decay |
| $\bar{\rho}, \sigma_{\rho}$ | 5, 1 | Normal distribution parameters of agent pushing forces |
| χ | 3ρ | Threshold of forced cell selection |
| ϕ | see below | Injury threshold |

6.4.2 Neighbourhood size

Our main interest is in evaluating the effects of introducing force into a microscopic crowd model. As discussed above, however, two changes were made to the basic model to facilitate

the addition of force: introduction of the η term into equation 9 and a switch to a four-cell neighbourhood.

We ran our implementation of the basic floor field model with and without these changes. The purpose was twofold: first, to ensure we replicated the basic floor field model from the FFM analysis, and second to determine whether the overall pattern of results was altered by our changes to equation 9 and the neighbourhood size. Results are shown in figure 7. On the first question, when running with the five-cell neighbourhood we found the model replicated the results of the FFM analysis by visual comparison with the FFM analysis (Kirchner & Schadschneider, 2002, figure 5b). On the second question, changes to the cell selection function and neighbourhood rules did not change functional forms of the results from the two




Figure 7: Non-force changes to the floor field model. Time steps to complete exit. Lines identified by \times are our replication of the FFM analysis results (Kirchner & Schadschneider, 2002, figure 5b). Lines identified by $+$ are our changed model (four-cell neighbourhood and η term).

models. The four-cell model did decrease the absolute time to complete exit, especially at lower values of k_S , but did not alter the overall relationships between k_S and k_D .

As the FFM analysis is not of a calibrated model, translations in the curves are acceptable for this baseline condition; we therefore conclude that our minor changes to the base model's rules are acceptable variations. We are less concerned with these practical changes and more interested in studying force effects. Accordingly, hereafter when we present results under the label of the FFM analysis we mean the floor field model running with our four-cell neighbourhood and the η term. In other words, the difference between the FFM analysis and our force analysis, below, is the difference between non-force and force-enabled models with otherwise identical neighbourhood and cell-selection rules.

6.5 Results: Effects of force

The simplest possible comparison with the floor field model is to compare it with the results of our model with injuries disabled. The most important results of the FFM analysis show the time for all agents to exit under the influence of varying k_S (Kirchner & Schadschneider, 2002, figure 5a) and varying k_D (*ibid.*, figure 5b). In this comparison we have set ϕ arbitrarily high, so that agents will not become injured but will otherwise experience the modelled force effects. We have chosen to present results for fixed k_S and varying k_D because these are the most revealing results. We have used the same values for k_S and range of k_D as in the FFM analysis. Results are shown in our figure 8.

The FFM analysis suggested that faster exit times are obtained when values of k_S increase, and that fixing any k_S value while increasing k_D results in a slightly longer time to exit. When k_S is low the force analysis suggests a similar pattern of exits. When k_S is high, however, the addition of force produces a dramatically different pattern of results. The number of iterations required to exit increases to between 402–1456% of the FFM analysis when $k_D \leq 5$ and $k_S = 10$. Where the FFM analysis indicated a clear benefit to high values of k_S , the force analysis suggests high costs for this parameter setting.




Figure 8: Time steps to complete exit with and without force. Lines identified by \times are the floor field model results. Lines identified by $+$ are with force enabled.

6.5.1 The effect of k_S on agent drive

To understand this discrepancy we must first understand agent drive, and the role of k_S in modulating it. The k_S parameter controls the agent's consideration of the static field (distance to the exit). A crowd-level effect that emerges from high k_S values in this scenario is that the area around the door becomes highly congested. The entire population of agents moves directly toward the door, leaving few to no empty spaces within the crowd. We describe the result of this movement independently of the size of the crowd in numbers or space by calling it a *high-density crowd* (figure 9a). Conversely, a *low-density crowd* has the quality of being spaced out with empty cells between individuals (figure 9b).

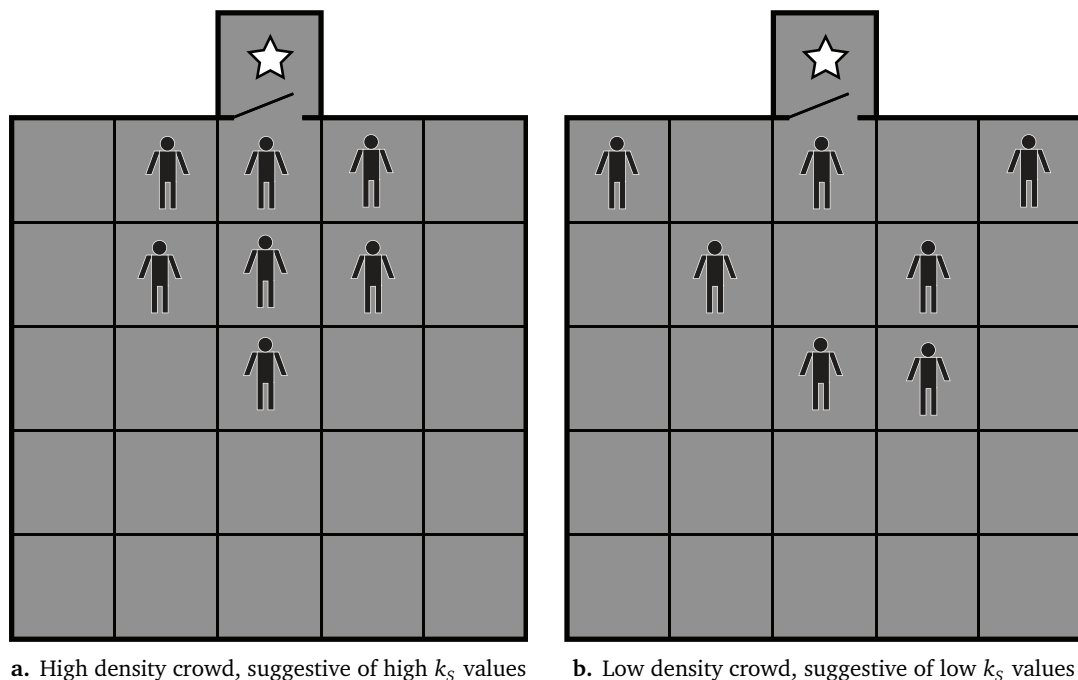


Figure 9: Crowd density schematic

Why does k_S produce this strong movement toward exits and high-density crowds? Table 3 demonstrates that when $k_S = 10$, it is virtually certain that an agent step will be in the direction of the closest exit. By contrast, when $k_S = 0.4$, the agent is more likely to step in a different direction than to step directly toward the door ($p_{door} = 0.3595$ vs. $p_{other} = 0.6405$). When $k_S = 0.4$ agents take much longer to reach the exit than when $k_S = 10$.

Table 3: Probability of selecting a neighbour cell while centred on the door, halfway back in the room. No neighbour cells occupied, no dynamic bosons in the neighbourhood.

| direction | k_S | | |
|-----------|--------|--------|------------------------|
| | 0.4 | 1.0 | 10.0 |
| forward | 0.3595 | 0.5378 | 0.9999 |
| left | 0.2394 | 0.1946 | 3.863×10^{-5} |
| right | 0.2394 | 0.1946 | 3.863×10^{-5} |
| backward | 0.1615 | 0.0727 | 2.060×10^{-9} |

Accordingly, we describe the strength of the k_S parameter as representing the agent's *drive to exit*. This is a key concept that we shall be returning to throughout the analyses of this chapter and those that follow. It is worth noting here that in the present scenario high drive and high density tend to occur together. However, the same is not necessarily true of all scenarios, as we shall see in future chapters. High drive is a consequence of the cell selection equations, while high density is a crowd-level effect that emerges from drive and the scenario.

Returning to the analysis of figure 8, we note that crowd density, generated by high drive, directly affects the effectiveness of force transmission. The differences between the models are minimal when the influence of the static field is low; when k_S (and consequently agent drive) is low there is little pressure for agents to move toward the door and agents tend to be dispersed, yielding a low density crowd. This low crowd density means lots of space around agents and therefore agents are infrequently frustrated in their desired movements by the presence of other agents; since agents must be blocked by other agents in order for force to be exerted, these low crowd densities translate into less force applied and therefore similar dynamics to the basic floor field model. In addition, the presence of empty cells means that when force is generated it is not easily transmitted throughout the crowd. This explains the similarity in results between the FFM and force analyses when $k_S = 0.4$ or $k_S = 1$. We now turn to the unique effects that emerge in the force analysis when $k_S = 10$.

6.5.2 Agent density and the aisle

Although our model (with injuries prevented) predicts similar exit times to the floor field model when k_S is low, it displays a dramatically different pattern when k_S (and consequently agent drive) is high. Times to exit increase significantly under these parameter values, particularly in combination with low k_D values. Under these conditions agents move directly toward the door at the outset of the simulation, creating a high density crowd. As k_S causes agents to prefer cells leading to the door, and as there are few open spaces within the crowd, agents tend to select occupied cells. With the selection of occupied cells, agents begin to push and this changes the dynamics of the system significantly from that of the floor field model.

This effect disappears as values of k_D increase because the importance of k_S is mathematically decreased (see section 5.3 for more on this point). This decreased influence of k_S results in lower agent drive, thus less co-ordinated movement toward the exit, lower density, more space around agents, less force applied and transmitted — and ultimately fewer deviations from the FFM analysis.

The longest exit times are obtained when $k_S=10$ and $k_D < 5$. Observation of the model shows characteristic patterns as the simulation unfolds with force on but injuries prevented (see figure 10). We shall first describe the pattern, and then discuss its mechanism. The agents pack themselves around the door at the outset of the model. There are few empty spaces within the crowd. The rear-most ranks of agents push toward the door. Forward of these agents, agents experience the combined forces of the agents behind them and lose their ability to select their own cells; this triggers additional involuntary force as their forced movements are blocked. Thus all agents join in a push toward the front, emanating from voluntary force from rear-most agents and compounded by involuntary force from those in front of them. Large additive forces are brought to bear upon agents at the front. This corresponds to the conditions shown in figure 10c. Our discussion now divides the agents into two groups: those in the area we call the *aisle* (a rectangular region, centred on the door, that extends front-to-back through the crowd), and those outside this region.

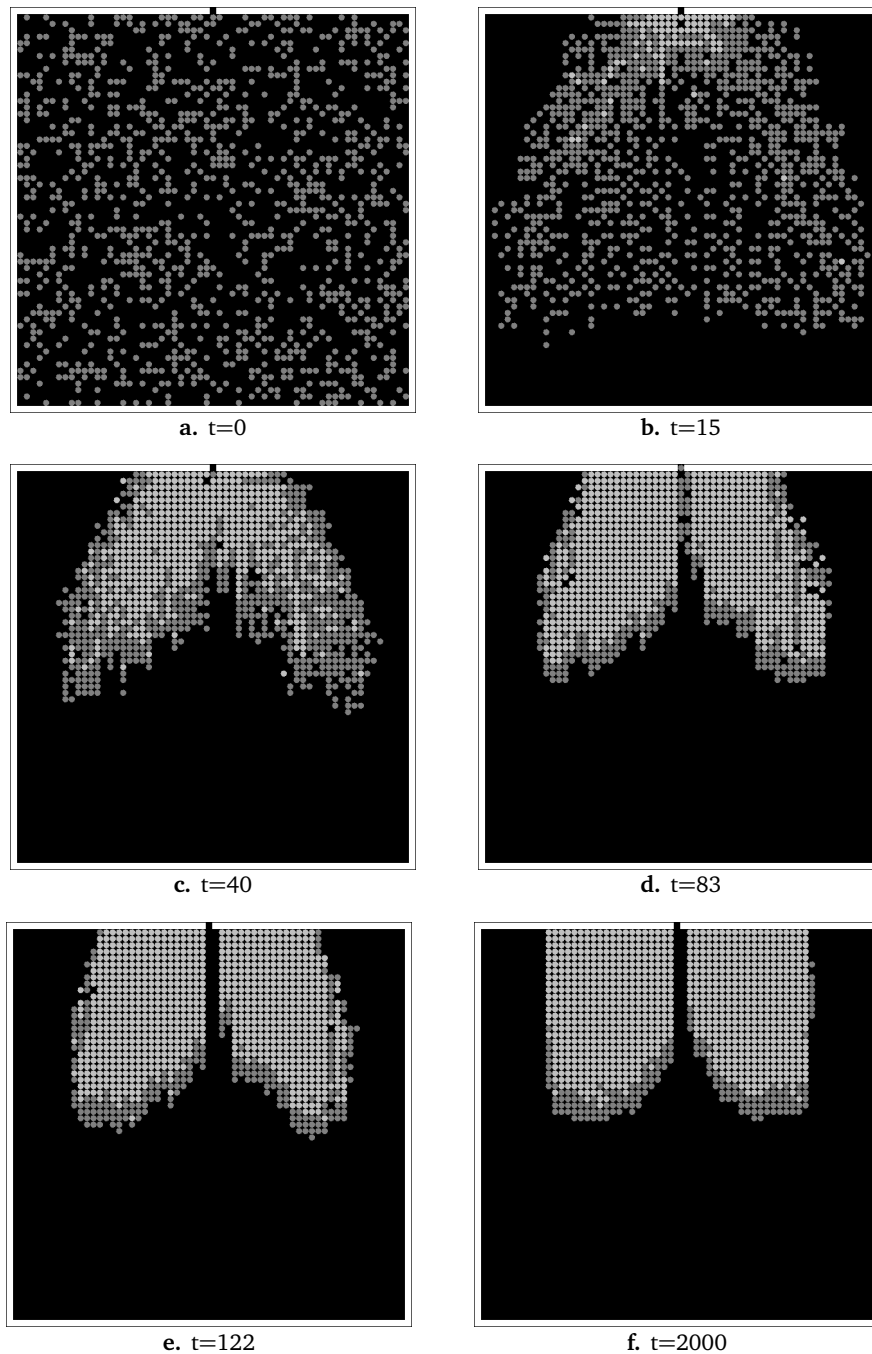


Figure 10: Aisle evolution. $k_s = 10$; $k_D = 3.5$; injuries off (a) Initial model conditions. Agents (dark grey circles) distributed throughout space delineated by walls (white border) must exit through door (black cell in wall). (b) time step 15: note agents from lower corners take longer to reach the door area because diagonal movement is not possible. (c) time step 40: Crowd around door now capable of efficiently transmitting force, and many agents (light grey circles) are now bypassing normal cell selection due to excessive force. (d) time step 83: Note free movement of agents on centre aisle. (e) time step 122: Aisle has exited, agents at edges of aisle and not at rear are light grey and cannot step sideways due to pressure from behind. (f) time step 2000: This stable pattern can persist for many time steps.

For agents on the aisle, the crowd pushes forward, and an agent at the front is forced to exit. This opens a space in the crowd that is immediately filled by an agent stepping forward. As time progresses, agents in the aisle are able to exit and the aisle gradually clears.

For agents not on the aisle, the rear ranks select cells on the closest path to the door (forward), applying voluntary force to those agents in front of them. In front of the rear ranks there is no opportunity for movement as agent cell selection is bypassed due to $\dot{f}_{ij} > \chi$. The involuntary forces from these agents are added to the voluntary forces from the rear ranks. The fully-packed crowd pushes forward upon the front agents who are pinned against the wall, and are thus prevented from exiting. In short, this demonstrates an emergent crowd-level effect in which non-aisle agents are effectively jammed in position due to pushing by a relatively small number of agents at the rear.

As the aisle clears (figure 10e), agents not in the aisle and not at the rear are physically prevented from stepping sideways into the empty aisle; this is because they are being forced to select the cell in front of them by the combined force of the agents standing behind them. This leads to a situation in which only the rear-most agents in the model are free to move. At this point, progress depends on one of the rear-most agents probabilistically moving laterally into the aisle, an eventuality that unfolds exceedingly slowly (see next subsection for more on this point). This process is laboriously repeated until enough agents have exited to alter the force characteristics such that agents close enough to the door (who prefer a lateral step to a forward step) are no longer pinned in place.

The distribution of results (increasing exit times) in figure 8 as k_D is increased from 0 to 3.5, is explained by differences in the initial packing of agents around the door. Higher values of k_D create a pack near the door that is shorter along the front wall and longer front-to-back. This results in a longer exit time because the exit takes place by slowly peeling off the ranks from the rear, and there are more ranks when k_D is higher. The value of k_D does not affect the formation of — or behaviour of — the aisle. As k_D exceeds 3, the decreasing exit times are due to the k_S parameter's reduction in influence as k_D becomes more important in cell selection (as discussed in section 5.3).

The results of this analysis support our claim that force effects are extremely important to studying crowd behaviours and hence for crowd models. The introduction of a simple force model of pushing without injuries results in dramatically different and more realistic exit behaviour in dense crowd conditions (see chapter 8 for more on this point).

Rear rank behaviour

The rear rank agents are the only agents in the model that can move after the aisle region is cleared. These agents overwhelmingly reject lateral or reverse movements, choosing to select the occupied cells immediately in front of themselves. In table 3 we showed the probability of stepping forward for an agent with no surrounding agents. Table 4 gives the probabilities for an agent blocked only from stepping forward (as the rear-most agents are in this case).

Table 4: Probability of selecting a neighbour cell while centred on the door, halfway back in the room. Only forward neighbour cell is occupied, no dynamic bosons in the neighbourhood.

| direction | k_S | | |
|-----------|--------|--------|------------------------|
| | 0.4 | 1.0 | 10.0 |
| forward | 0.2191 | 0.3678 | 0.9998 |
| left | 0.2919 | 0.2663 | 7.726×10^{-5} |
| right | 0.2919 | 0.2663 | 7.726×10^{-5} |
| backward | 0.1969 | 0.0995 | 4.121×10^{-9} |

As discussed in footnote 1 on page 85, polarisation in neighbourhood cell scores can underlie increased probability that agents will select occupied cells. When $k_S = 10$, this results in doubling of the probabilities of selecting the non-forward cell, but this doubling is insignificant compared to the overwhelming probability of moving ahead. (Compare with $k_S = 0.4$, where moving either left or right exceeds the forward movement probability in this case.) While table 4 assumes only one blocking agent, for most of the rearmost agents several of the neighbour cells may be occupied, further diluting any inhibition for forward movement. This is the explanation for the persistence of the aisle region as shown in figure 10f, where rear-most agents are very slow to move laterally into the aisle.

6.5.3 The aisle and model realism

The fixed aisle pattern appears to run counter to the normal behaviour we expect from crowds; it is hard to believe that, in reality, people at the rear would not simply step sideways into the aisle.³ In this regard, it must be highlighted that we have disabled part of the model by setting $\phi = \infty$. In other words, the aisle forms (a) with high agent drive, and, importantly (b) *when we prevent injuries*, but infinite force with no injuries is no more realistic than rear agents who do not make lateral steps. As we shall see, injuries play a key role in improving aisle breakdown (because this long-lived aisle pattern is obtained when the agents have an infinite capacity to experience force without becoming injured). Injuries, however, can block the exit, preventing a run to completion and consequently preventing comparison between the force analysis and the FFM analysis. Hence the change in measure to be studied in the next section, and our reasons for setting $\phi = \infty$. We shall have more to say on the topic of the aisle in section 8.4.

One remaining question is whether $k_s = 10$ — given its effect on movement probabilities — is a useful parameter value for study. We must bear in mind that we are working with an abstract model, and we do not know the range of realistic values for this parameter. It is important to note that, although the $k_s = 10$ parameter value produces almost deterministic behaviour, it also generates a scenario we want to study in this chapter. Here, our goal is to evaluate the effects of pushing — including extreme force build-up — and the forward focus of agents is generating this interesting force build-up. We have seen cases (e.g. Bethnal Green, Riverfront Coliseum) in which those who have freedom of action at the rear push and create shockwaves — unaware of problems engendered at the front — just as our rear agents do. Fruin, moreover, discusses the loss of individual control in tightly packed crowds, and describes crowd deaths in terms of the kind of force build-up we see in this model. Replication of these conditions is a valuable aspect of setting $k_s = 10$, as is comparison with the FFM analysis. We also note that section 6.4.2 and figure 7 showed good agreement between the FFM analysis and our basic model (with the

³If we take this position at face value, the interdiction of diagonal movements constrains agents in this regard. This is because a choice must be made between moving sideways and moving forward. If simultaneous forward-diagonal movement was allowed, it would increase the probability of rear agents stepping into the aisle region (perhaps effectively preventing the formation of the aisle in the first place). This observation will motivate a suggestion for future work: the addition of diagonal movement choices to the model. The question, however, is academic because the lack of injuries in the present scenario means that the appearance of the aisle is artificial in any case.

non-force changes of section 6.3.2); thus it is the introduction of force alone that is responsible for the aisle region. In the interests of investigating high forces, then, we conclude that $k_S = 10$ has value for present purposes; in a study less interested in extreme force conditions, however, a more modest value may be appropriate.

6.6 Results: Comparison of measures

As discussed above, it is necessary when analysing injuries to use a measure that does not require all the agents to exit; this is due to the potential for the door to become blocked by an injury. We ran the model, with and without the force simulation, using a *fixed time* measure. We ran the model for 350 iterations and measured the number of agents exiting within that time. Results are shown in figure 11. The value 350 was selected because it allows for the initial stages of model exit to be completed, it excludes most of the aisle iterations and it precludes a complete exit for the fastest runs (which would lead to a floor effect obscuring differences between fast runs). Note that to bridge between the complete exit and fixed time measures, we again evaluate the model with injuries prevented ($\phi = \infty$).

In inspecting graphs of the present results with respect to the ones just described, bear in mind the reversal in sense of the y -axis: In the complete exit measure, a low number of time steps is good: a fast exit; this corresponds in the fixed time measure to a high number: large numbers of agents exiting in the fixed time.

There are broad similarities in the pattern of results between the complete exit measure and the fixed time measure. As in figure 8, when $k_S = 0.4$ and $k_S = 1.0$ we see a similar pattern of results between the FFM and force analyses. For both models, both measures and both values of k_S , increasing k_D indicates less efficient egress (increased time to exit in figure 8, decreased agents exiting in figure 11). We see an improvement in exit times, especially when k_D is low, when force is applied compared to the non force cases; this effect is more pronounced when $k_S = 1$, using both measures. As before, the FFM analysis predicts maximal exit rates when k_S is 10, while the force analysis disagrees, showing many fewer agent exits within 350 iterations, especially when k_D is low.




Figure 11: Number exiting in fixed time with and without force. Lines identified by \times are the floor field model results. Lines identified by $+$ are with force enabled.

This demonstrates two points: First, the outcome of the complete exit is visible from the crowd's progress in the early stages of the model's execution; we do not need to wait for a complete exit to compare the effects of the various parameter values. Second, the new fixed time measure detects similar trends and effects when compared to the complete exit measure. This will allow us to relate our exploration of modelled injuries (below) with the FFM analysis despite the difference in measures.

There is one notable difference between the two measures, however, concerning the results when $k_D < 5$ and $k_S = 10$ with force. In the fixed time measure the curve is smooth and does not experience a drop in exiting agents centred on $k_S = 3.5$, as was apparent in figure 8. This is a reflection of the timing of the exits with force enabled. The fixed time measure does not explore the slow exit phase of the non-aisle agents.

6.7 Results: Effects of injuries

With the fixed time measure in place, we can enable the injury parameters of the model in order to assess the effects of injuries on exit behaviours of agents. In order to cause agent injuries it is necessary to determine another threshold value in the agent's injury response. When an agent experiences force greater than the threshold, ϕ , they are deemed to be injured. The ϕ parameter represents hardiness; the larger the ϕ parameter, the more force the agents can experience before becoming injured. Like other aspects of the floor field model, we have yet to determine the correspondence between model units of time, space and force and the physical units experienced by real crowds. Rather than selecting a threshold arbitrarily, we have elected to characterise this particular variable across a range of values. (The performance of a particular individual is hypothesised, then, to lie at some point along this axis.)

In gathering model data, we have selected representative values for k_S and k_D from the FFM analysis. We noted that Kirchner and Schadschneider selected fixed k_D values of 0 and 10 (2002, figure 5a) and k_S values of 0.4, 1 and 10 (*ibid.*, figure 5b). We matched each k_D value with its opposite extreme from the k_S values. In creating a matching value for the intermediate $k_S = 1$ we selected $k_D = 4$, which occurs approximately halfway along the range of complete exit times

for $k_S = 1$ in Kirchner and Schadschneider's figure 5a. This gives us a set of values (k_S, k_D) : $\{(0.4, 10), (1, 4), (10, 0)\}$. By analogy with the discussion of section 6.5.1, we refer to these as *low drive*, *moderate drive*, and *high drive* cases, respectively (and we shall use these terms in future chapters as well). For each of these parameter regimes the injury threshold was varied between 25 and 225 bosons, exploring the injury space up to a point where the threshold was too high for any agent to become injured. We report numbers of injuries and number of agents exiting in a fixed time in figure 12.

Unsurprisingly, as increasing ϕ required more force to cause injury, injuries decreased. In terms of numbers of injuries, as drive increases (for a fixed injury threshold) more agents become injured.

Turning to the exit chart, when $\phi > 155$, injuries do not occur and we see familiar values that correspond with the points in figure 11. Introduction of injuries affected all drive levels, but primarily the moderate and high drive cases, with exit rates affected by changes in ϕ . The parameter settings that were the best for preventing injuries were different from the parameter settings that enabled more agents to exit. The best exit rate was produced in the medium drive case (and this accords with the results of sections 6.5 and 6.6 as well). A lower exit rate was obtained from the low drive case. A non-monotonic relationship was found for the high drive case: at high and low ϕ , the number of agents exiting was low; for $45 < \phi < 135$, however, exit rates were substantially higher with a peak at $\phi = 95$.

These general trends and relationships between drive levels are robust, and have been observed consistently by us over many replications of these data sets. Caution is urged, however, about attempting prediction of the results of an individual run from this data; distributions for number of exits at values of ϕ that produced injuries have high standard deviations (see figure 13). These high standard deviations are obtained because of the possibility of an injured agent (or group of injured agents) blocking the exit.




Figure 12: Number injured and exited in fixed time with force and injuries. Upper results for injury, lower results for exits.




Figure 13: Standard deviation of number exiting in figure 12.

6.7.1 Mechanism underlying injury effects

Now let us consider the mechanisms that underlie the results of figure 12. In the terms of section 6.5.1, agent drive — and hence crowd density — has a dominating effect on both injury production and numbers of agent exits. High drive where $k_S \gg k_D$ produces more injuries, while moderate drive when k_S and k_D are both moderate produce the best numbers of exits. Low crowd drive brought on by $k_S \ll k_D$ results in few injuries but also few exits.

At low drive the lower number of injuries is a simple consequence of the lack of force exerted in this case, as described in section 6.5.1. The two other conditions are more revealing.

In the moderate drive case there are three separate phases displayed. The first phase ($\phi < 75$) sees a substantial drop-off in injuries accompanied by a substantial increase in the number of agents exiting. The second phase ($75 < \phi < 145$) sees a continuing decline in injuries until no agents are injured, but only a modest increase in number of agents exiting. The third phase ($\phi > 145$) involves a consistent lack of injuries and no change to the exit rate. This

indicates that a moderate number of injuries does not necessarily have a large effect on the number of agents exiting when drive is moderate. Larger numbers of injured agents do have an impact on number of exits.

When $k_s = 10$ the cell selection policy provides a strong pressure to move toward the door, and consequently high density near the door, as described in section 6.5.1. The ϕ parameter seems to have a simple relationship with the number of agents injured: we see a decline in numbers of injuries as ϕ increases. This simplicity in *numbers* of injuries, however, belies a complexity in the *effect* of injuries on the exit rate, which display a complex relationship against ϕ . There are four phases: First, when $\phi \leq 55$, increasing force threshold results in an increase in exits to a local maximum. Second, while $55 < \phi \leq 95$ a brief plateau occurs, then exits build to a maximum. Third, when $\phi > 95$, number of exits decline, until — fourth, there are no more injuries and the exits level off.

Protective effect of injuries

The interpretation of this four-phased relationship at high drive is made simpler by the fact that obstacles can actually facilitate exits in emergencies. This is because obstacles disrupt symmetrical patterns of force and allow shelter for exiting agents from the forces behind them (Helbing et al., 2002; Nishinari, in press). This effect is evident in our model, and is the underlying mechanism for the increased exit rates between $45 < \phi < 145$. This swell in agent exits is attributable to the protective effect of the injured agents upon the others. Practically speaking, because force is not transmitted through injured agents, those just forward of the injured are free to make their own cell selections while others are pinned in place; this increased flexibility means that agent exits are facilitated, and avoids the kind of stasis evident in figure 10e, where drive and density were high but we artificially prevented injuries. Observation of the model demonstrates this principle at work in several different ways. Let us take two examples:

At $\phi = 55$ large numbers of injuries can create zones of free movement (see figure 14). These large numbers of injuries, however, physically remove many agents from the simulation (directly limiting exits) and impede movement for those behind (indirectly limiting exits).




Figure 14: Large force block with many injuries. $\phi = 55$, $k_S = 10$, $k_D = 0$, time step 120. Dark grey agents are making their own cell selections, both at the rear, and in front of injured agents (yellow). Light grey agents are pinned in place.

The highest number of exits, around $\phi = 95$, are more subtle. They take advantage of limited numbers of injuries to perturb the aisle pattern. In figure 15, an agent near the front (agent a) is shown stepping laterally into the aisle, an event that would not occur with injuries disabled. This step is possible because to the right of this agent there is an agent (agent b) sheltered by injuries from other agents' pushing and leaning forces. Agent b is therefore capable of selecting its own cell. Although b 's force exertion is small relative to the many agents pushing upon a from the rear, through re-propagation by the row of agents between a and b it is able to shift the force upon agent a slightly to the left. Equation 8 then provides a small probability for a to step into the aisle. This movement opens a space, breaking the chain of force transmission and allowing a column of agents to break away into the aisle.

6.8 General discussion

Having explored the characteristics of the force introduced by our changes, we now turn to the implications of this exploration for our two central issues: the predictions made by the FFM analysis and the centrality of force to models of crowds.

Exit rate. First we consider the overall relationship between the two analyses. Table 5 summarises the number of agents exiting for the values of k_S and k_D investigated, reporting the

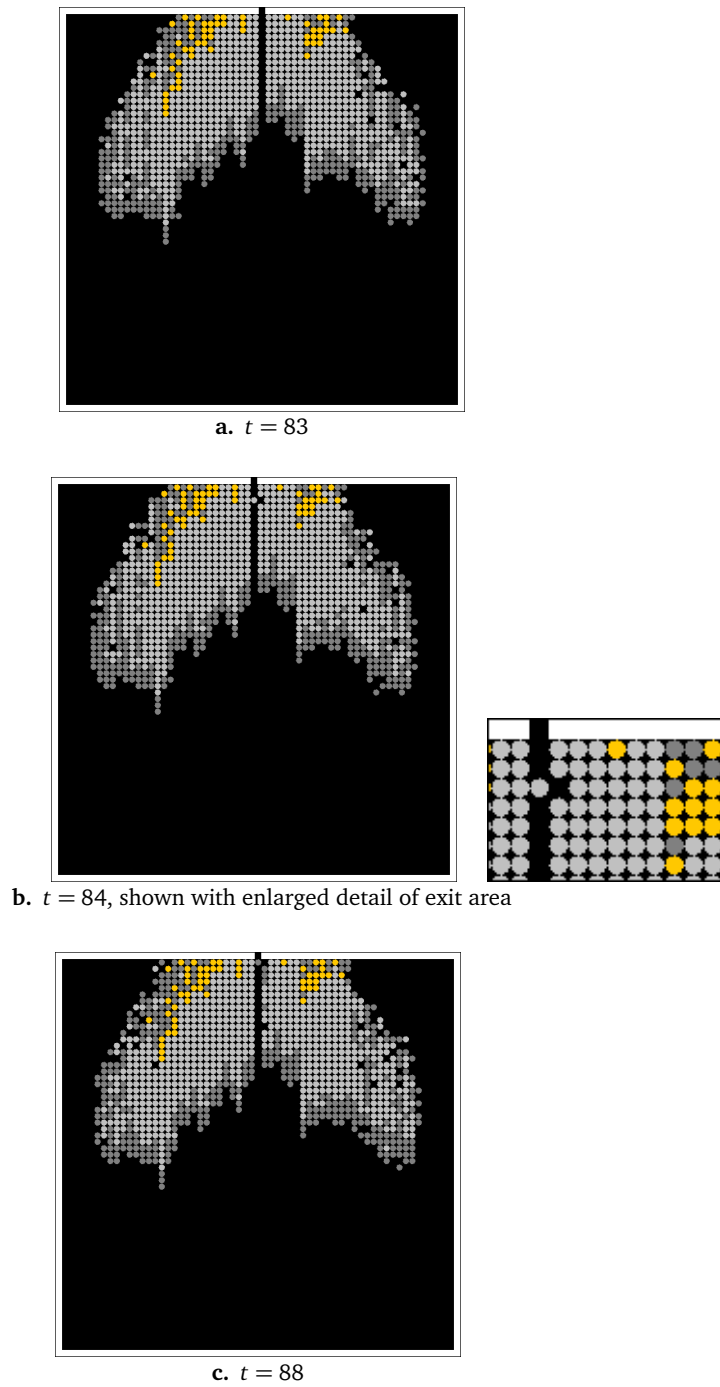


Figure 15: Devolution of the aisle. $k_S = 10$, $k_D = 0$, $\phi = 105$. Note exiting agent at $t = 84$ stepping left into the aisle.

Table 5: Comparison of FFM and force analyses (agents exiting). Differences in quantities identified by the same symbol are statistically significant ($p < 0.0001$, two-tailed welch’s t-test).

| drive | (k_S, k_D) | FFM analysis | Force analysis | |
|----------|--------------|--------------------|----------------|-------------------------|
| | | mean | mean | range |
| low | (0.4, 10) | 55.9 | 54.5 | 40.0–57.4 |
| moderate | (1, 4) | 118.0 [♦] | 92.0 | 36.3–106.1 [♦] |
| high | (10, 0) | 174.3 [▲] | 50.8 | 16.7–85.7 [▲] |

mean and range across ϕ for each crowd drive.⁴ Force has a significant impact on the number of agents able to exit in the first 350 time steps of the simulation, unless agent drive is low (conservatively, we compared the best exit rate in the range of the force analysis to the FFM analysis). In the high and moderate drive cases the range of agents exiting is below the number of agents exiting in the floor field simulation; reductions of 50–90% in exits (depending on the threshold for injury) were obtained at high drive when the force analysis predicts injuries in large numbers.

On the question of facilitating agent exits, table 5 demonstrates that when agent pushing is taken into account, our model makes a very different prediction from the floor field model. Specifically, our model suggests that conditions that reduce crowd drive to moderate levels (namely moderate values of k_S and k_D) will result in many more agent exits than will high crowd drives (as predicted by Kirchner & Schadschneider, 2002).

Crowd paradoxes. Helbing, Farkas & Vicsek have observed that crowds exhibit paradoxical collective behaviours when speeds and commitment to courses of action increase. These behaviours are characterised by simple paradoxical phrases like faster is slower and freezing by heating (see section 2.1.2 for more on these paradoxes). The force analysis demonstrates these two paradoxes.

⁴Table 5 presents mean data for the force analysis in order to provide the reader with an overall sense of the effects of injuries caused at each parameter level. From an analytic perspective, however, an average across values of ϕ may not be warranted as we hypothesise that there is a single true value of ϕ for each individual. We prefer to analyse the range data as this provides a best-case and worst-case prediction of the effects of injuries on the exiting crowd.

The faster is slower paradox (Helbing et al., 2000b) refers to the fact that crowds passing through a bottleneck may make slower progress when individuals' speeds are higher, due to the enhancement of clogging at the exit and interpersonal force effects. Our model directly demonstrates this principle as moderate speeds toward the door (brought on by moderate parameter values) yield the best exit rates. Faster speeds (brought on by $k_S \gg k_D$) yield fewer agent exits in a fixed interval because of clogging and also force effects.

The freezing by heating paradox (Helbing, Farkas, & Vicsek, 2000a) concerns particle energies: High particle energy can result in stasis where lower energies permit free movement. Our model directly demonstrates this principle: High particle energy (in the case of pushy, dense crowds) demonstrates the stasis effect of figure 10e that is only broken when the energies are dissipated through force breaks — a demonstration of freezing by heating.

Is force essential? This brings us to a matter raised at the start of the chapter in which we suggested that there were three possible positions that a model could adopt on the question of crowd forces.

The first position suggested that modelling force was non-essential and that results could be accurate with or without force. Based on our investigation of the same model both with and without a force implementation we believe that this position cannot be supported. The addition of force to the model, with or without injuries, dramatically and significantly alters the exit behaviours of modelled agents.

The second position suggested is that force can be treated *post facto* by applying some correction factor to the data. The complex, multi-phase relationships of figure 12, and our understanding of this relationship by reference to the protective effects of injuries at high drive, suggest that the effects of force cannot be modelled after the fact by applying a simple “force penalty” measure if a range of crowd conditions modelled is to include high drive crowds. No force penalty measure can induce the non-linear changes we have reported.

Instead, our results support the third conclusion, that force is characteristic of crowds and cannot be abstracted from a crowd model without significantly changing the behaviour of the model.

This conclusion demonstrates the value of studying human behaviour in a microscopic crowd model. The abstract simulation allows for us to test general hypotheses such as this without becoming distracted by the details of specific situations. But what of the validity of the simulation? We now turn to a discussion of how the model relates to the world.

6.9 Model validity

As discussed in section 4.3.2, the Microscopic Human Factors methodology gives us the opportunity to evaluate the model at all three levels of abstraction. In this section we argue for the validity of pushing behaviours in the floor field model.

Implementation. At the level of the implementation, the task is to determine whether the model has properly implemented the microscopic human factors that were specified in the reduction (section 6.2). The reduction specified that force is generated purposefully, directed, propagated and carries consequences. Through our analysis of the floor field model we have shown all of these properties at work within the model. Purposeful force generation is demonstrated as agents push primarily when crowds are dense and agent movements are blocked. Directedness and propagation have been demonstrated as forces are directed toward the door, pinning agents in place as we have seen in figure 10f. Our injury investigation has demonstrated non-linear consequences of force. As discussed above, the implementation may benefit from the addition of diagonal movement transitions, a possibility that we leave to future work.

Reduction. The validation question at the reduction level concerns whether the microscopic human factors proposed are a reasonable abstraction of our conception of how pushing works in real crowds. The behaviours generated by our microscopic human factors do seem to accord with the specification. The specification was based primarily on Fruin's observations about force in crowds. Our reduced rules generated the pushing and leaning forces he speaks of. Observation of the model shows waves of force that propagate through the crowd (figure 16), particularly when crowds reach maximum density.




Figure 16: Waves of force. $\phi = 105$, $k_S = 10$, $k_D = 0$, time step 199. Lightness of cell proportional to number of force bosons present. Note wave pattern in intensity variations toward edges of crowd. Agents not shown.

Consideration of the reduction also leads to potential future work. Our reduction of the specification simplified the concept of injury in crowds to immobility and physical obstruction. Injuries in crowds, however, have a continuum of severity from moderate discomfort to death. Clearly our abstraction in this case has discarded this detail. For our present purposes this allowed for a clear focus on the role of injuries as force breaks. Future work could integrate more subtle injury models with a conception, for example, of variable speed (e.g. Kretz, 2007) to examine the effect of lightly injured mobile agents on egress dynamics.

The results we reported in previous sections show that our reduced rules are capable of re-creating pushing, leaning and simulation of forces in a crowd situation. In our model, people have individual control only when drives and densities are low and forces are controlled. We have demonstrated a crowd in which movement can be involuntary, and in which force propagates. We have provided for simple simulated injuries. In short, our reduction meets the goals of the specification.

Specification. Validation at the level of the specification means determining the extent to which the model results are relevant to the real world. As discussed in section 4.3.2, the methodology helps us to argue for sufficiency of the model results, helping our results to fit in with the results of other disciplines to confirm the mechanisms at work in crowds.

We have shown that introduction of voluntary pushing based on an individual's frustration in movement is capable of producing crowd behaviours that accord with Fruin's description of forces in crowds when combined with simulated leaning forces. We have shown emergent effects reported by Fruin (e.g. shock waves) in model results. We can look to all three case studies in chapter 2 for additional support in our contention of the importance of modelling forces in crowds. All three case studies showed a profound effect on movement due to high drives, high density and force: injuries causing outright blockages at Bethnal Green tube station and Lenin stadium, and force causing injuries, loss of control and pinning of agents at walls at Riverfront Coliseum. These sources of support suggest validity for our model, in turn supporting the conclusions of our work relating to the importance of modelling force in crowds and the protective effect of injuries in their role as force-breaks within the crowd.

As discussed above, however, questions concerning the exact trigger for pushing behaviours are not answered by the model. The model, then, should therefore not be seen as a complete model of crowd safety. Our abstract model, instead, is useful to characterise mechanisms for emergent crowd effects due to pushing and force. We show what happens when force is deployed, and show a mechanism that can generate the force required to conduct this analysis.

6.10 Summary

In providing the first of three demonstrations of the MHF methodology, we created a model of pushing in crowds, force and crowd safety. We used the MHF methodology to describe the behaviour of interest at the specification, reduction and implementation level. We then used the implementation to examine the mechanisms for force propagation, injuries and related effects at the crowd level. We specified what parts of the physics of force need to be translated into a model, and showed how the addition of crowd safety alters model conditions from an established microscopic model of crowds: the floor field model.

Conclusions from this work include the fact that a relatively small number of agents in the rear of a crowd can dramatically affect outcomes for those closer to the front, including jamming effects due to leaning forces throughout the crowd triggered by the rear agents. We have also

noted that relatively smooth injury rates belie complex relationships between injuries, crowd drive and number of agents able to exit: at low drive injuries directly prevent exits; at moderate drive a moderate number of injuries does not affect the exit rate; at high drive, as injuries provide a protective effect, they can act as force breaks in facilitating exits. Contrary to the results of the floor field analysis, our results suggest moderate agent drive yields the best exit performance, and suggests that reducing high drive would increase safety and exit rates. This is a key result that will be used in other demonstrations; we shall also return to this point in the general discussion in section 9.1.2.

By demonstrating force-related non-linear changes in numbers of agents exiting we have shown that force must be treated as an essential aspect of crowd models. Important effects like the facilitative aspect of injuries are lost if force is not considered, or if it is treated only as an add-on parameter that scales the output of a model that otherwise does not consider force effects. We have shown a much more complex pattern of exit times as compared with the floor field analysis — both with the complete exit measure adopted by Kirchner and Schadschneider (2002), and also using a fixed time measure that enables investigation of injuries (that could otherwise block the door).

We have discussed the model's validity, and shown how the methodology can argue for a model's appropriateness, while at the same time discovering aspects of the implementation and reduction that need further investigation, and identifying areas of uncertainty in the literature. Future work to improve the simulation of force could include provision for diagonal movement, and a more refined conception of injuries. Although Fruin's remarks indicate that force emanates from individual pushing, little work exists in the literature to identify conditions that favour individual expression of pushing behaviours; obtaining new facts in this area would significantly improve models of crowd safety. Additionally, relation to real-world parameters like distance, time, force thresholds and agent strength could be of interest, always bearing in mind the focus on abstract situations. Additional experimentation could also be done with the present model; it should be possible to use the model to suggest patterns of wall placement that can act as force breaks: facilitating flow, reducing agent pinning and injuries, and ultimately increasing the exit rate.

Returning to the purpose of the three demonstrations raised in the opening paragraph of this chapter, we identified three main goals: demonstrate the practicality of the methodology, demonstrate the ability of the methodology to create models that take account of Sime's concerns, and to show that the models created have value in generating discussion and conclusions. On the first point, this chapter demonstrates the practicality of using the methodology to generate not only an implementation of a human behaviour (voluntary pushing) but also to generate an implementation of an individual-based physical simulation (force and crowd safety). On the third point, as just described, the model has generated an understanding of the nature of forces in crowd safety, a fact that we take to justify our contention that the methodology can generate useful results and discussion directly from the models created. The second point, of addressing Sime's concerns, bears some additional consideration: The demonstration adds a very simple behaviour (voluntary pushing) triggered by agent frustration in movement, and has added some richness to the human behaviour simulation in the model. The behaviour, however, is undeniably quite simple and we have spent as much time in this demonstration on simulating force itself (an investment that will be well repaid in the following two demonstrations). The emergent force-based effects at the crowd level are physical in nature, but emanate from an individual's decision to push, and are thus legitimately emergent crowd-level effects flowing from human behaviour in the form of voluntary pushing. Thus, we have come some of the way to adopting Sime's concerns with this demonstration.

The goal of our second MHF demonstration, in the next chapter, will be to add a much more overtly knowledge-based behaviour simulation to agents, and will represent a much more direct demonstration of the integration of Sime's concerns with a microscopic model. We will build on the results of the present chapter to evaluate how differing goals between agents affect agent exits and safety.

Chapter 7

Discovery and communication of spatial information

The second demonstration of the methodology is to directly cause heterogeneous behaviour, and to use the microscopic model to investigate the interactions between the groups created. The simplest way to implement a microscopic model is for all agents to be identical, processing the same information to the same ends: a homogeneous crowd. Real crowds, however, are heterogeneous; different people in the crowd move according to differing goals and change their behaviour based on local conditions (e.g. Johnson, 1987b; Still, 2000). Extreme cases such as the Lenin stadium disaster (section 2.2.2) highlight the fact that multiple goals within a single crowd can cause radical departures from normal crowd behaviour. When models use a homogeneous approach, they limit their applicability to these kinds of situations.

As discussed in section 4.3.1, the creation of any crowd model rests on a conception of how the modelled phenomenon occurs in real crowds. In contrast to the force model of chapter 6 and the front-to-back communication model of chapter 8, the purpose of modelling differing points of view is not the investigation of a specific individual behaviour to understand the emergence of a group property. Rather, we are interested in using a microscopic model to look at the effects of an emergent group property (heterogeneity) itself. There may be many different factors that underlie heterogeneity in real crowds, and it is not necessary for us to identify all of these factors in order to produce a heterogeneous model. We have chosen to focus on discovery and communication of spatial information because differences in spatial knowledge will underlie different patterns of movement and therefore a heterogeneous crowd. In chapter 8

we shall continue our interest in information by investigating a different type of communication: front-to-back communication.

7.1 Specification

In adding discovery and communication of spatial information to the model we have been inspired by Jonathan Sime's observation that:

It is important to note that a building or setting (such as an underground station) is not only a physical space or structure, but *an information system through which people move*. If this is remembered, it should help direct attention to the perspective of crowd members, as well as to the physical dimensions of a setting (i.e. psychological and engineering parameters). (1995, emphasis Sime's)

Just as our specification of force was guided by Fruin's observations of the effects of force in crowds, our implementation of discovery and communication is inspired by Sime's comments. Sime emphasises physical aspects of the space and the perspective of crowd members (psychological and engineering views) but also a third concept: that of an information system through which people move. An information *system* that can be *moved through* implies that the system has components that are distributed through space and that are both perceptible and informative. Informative means "having the quality of imparting knowledge or communicating information; instructive" (*Oxford English Dictionary*, 1989). Thus, Sime is describing a system in which people gain information by moving through a space in which various locations afford this instruction.¹ To a person with more complete knowledge, the space affords more options in terms of movement goals.

That people can extract information from their environment implies that two people close to one another may be engaged in different actions. In exiting a space, a more knowledgeable person may move differently (e.g. toward little-known exits) from others close to them, a difference explained by their extra knowledge. People moving through a space are sensitive to new information and can readily absorb it and change their behaviour accordingly. For example,

¹This locality of some information, however, does not necessarily preclude global information; overhead announcements may occur during an egress situation that allow people spread over a larger area to accumulate knowledge of the space as well.

upon approaching a blocked exit, a person wishing to leave changes their behaviour (e.g. moving instead to an alternate exit). Similarly, a person discovering a previously unknown open exit will assimilate this information leading to the potential for its use. (It should be noted that this premise, in particular during an evacuation, is contradictory to the hypothesis that people panic. It depends on a view of people as rational and capable of cognitive processing.)

It is important to note that it is not only inanimate aspects of a scenario that can carry information. People moving through a space occupied by others are capable of sharing their knowledge of the space through communication. Johnson has reported that shouting was futile in a crowd disaster (1987b), so we adopt the conservative view that effective communication between crowd members (as distinct from, say, amplified overhead announcements by crowd managers) must be direct and personal.

It must be acknowledged that there is much that is known about information gathering and heterogeneous crowds that is not captured by this specification. For example, issues of trust in communicated information and the desire to use familiar exits are not captured by the premises we have adopted. As discussed in sections 4.2 and 4.3, the microscopic human factors approach is to simplify and abstract human behaviours in order to bring them into a microscopic model at an appropriate scale. For the purpose of generating a heterogeneous model we do not need to bring in additional complicating factors, and we leave their investigation to future work.

7.2 Reduction

Having explained our conception of knowledge-based heterogeneity in crowds, we turn now to a discussion of how we can provide for these mechanisms as microscopic human factors. The goal in reduction is to develop an abstract, formalism-neutral, individual-focused model of the behaviour that can account for the specification. With this in mind, we aim to model heterogeneity by introducing a simulation of divergent movement goals within the modelled crowd. These divergent goals will arise from evolution of the beliefs of individuals, caused by discovery and communication of new spatial facts. The ultimate source of these facts will be an information system incorporated at the heart of the movement model. Unlike the previous

demonstration, all of the modifications to the model will be toward the implementation of a simple cognitive behaviour simulation for crowd members. From the specification above, we derive the following essential, and individual rules:

Explicit mental content. The key to generating a heterogeneous model based on discovery and communication of knowledge of space is to note that different people have differing mental content about their environment. In a microscopic model we aim for simplicity and abstraction; we want to avoid complicating the model with, for example, inductive or deductive reasoning. In order to determine the effects of — and interactions between — different beliefs on behaviour, it is essential that the set of possible beliefs be restricted in size. Thus, our abstract model maintains certain specified and indivisible views of the world, and agents select from among these views during model execution. These views are created by the modeller as part of specifying the scenario.

Discovery. People accumulate information about their surroundings as they move throughout the space. This discovery is triggered upon moving to particular circumscribed areas within the model, whose size and placement is consistent with the modeller's intentions with respect to the perceptual capabilities of crowd members being modelled. Discovery of information corresponds to changing mental content as described above. Like the mental content itself, the discovery areas are specified by the modeller, who may choose to allow these discovery zones to be somewhat dynamic; this allows representation of information that becomes available at certain locations after the start of the simulation (e.g. localised announcements).

Accumulation of knowledge. Knowledge of an individual in a crowd is accumulated over time. Learning some facts about the world does not preclude additional facts being learned later, and later learning generally accumulates with — as opposed to supplanting — previous mental content. To some extent this rule is already addressed by the set of indivisible mental content, which requires the modeller to specify all possible views of the world. The present rule, however, goes further by requiring the provision of acceptable transition rules. For example, a

person discovering that an exit is blocked will not generally regress to a state in which they have forgotten the blockage. (This is not to say that misinformation cannot occur. Someone can learn of a putative exit and subsequently learn that the exit does not exist. In this case it is not that they have regressed to an earlier state of knowledge through simple forgetting.)

Directed communication. By simulating direct communication we can use the model to investigate behaviour engendered by people becoming informed, including misinformed, without physically visiting the discovery locations. As noted in the specification, the communication model is assumed to rely on direct person-to-person contact. As with the force simulation in chapter 6, we require a simple, abstract trigger for communication. We take communication to occur not on a random basis, but with a purpose, namely to assist in achieving movement goals. In this sense, communication has a similarity with voluntary pushing; both are an attempt to get a conflicting agent to adopt a more complementary movement pattern. We assume that communication, therefore, will occur exactly in the same situations in which the agent applies a pushing force: when movement is frustrated due to an occupying agent on the agent's desired cell. (Other possible purposes for communication, for example local leadership and altruism, are left for future work.)

7.3 Implementation

The MHF methodology requires that this reduction be implemented with close fusion of the reduced rules and the existing rules of the floor field model. The implementation discussed in this section is presented in pseudo-code in algorithm 3 on page 126.

The changes to the model depart from the floor field model in two major respects. First, they allow for individual agents to perceive the modelled world differently from one another, creating a heterogeneous crowd. Second, they allow for agents to change their view of the world, either under the influence of a new discovery field, or through a simple inter-agent communication mechanism.

7.3.1 Representations of space

In the floor field model, the representation of space is the static field, which encodes a distance from points of interest. The reduction requires that the modeller make explicit all the possible views of the world. In order to represent multiple views of the world within the floor field model, we replace its single static field with a set of static fields. Each possible view of the world is represented as a distinct field within this set, and each represents an indivisible and complete view of the space.

We redefine the symbol S to refer to a set of static fields rather than the sole static field of the floor field model:

$$S = \{S_0, S_1, S_2, \dots, S_n\} \quad (10)$$

The semantics of the static fields are not altered from the original model. Each of the fields is created from a set of points of interest using one of several metrics (Nishinari et al., 2004; Kretz et al., in press) to determine the distance from each cell to the nearest such point. The last chapter took the set of locations of interest to be the same as the set of exit cells in creation of a single static field; in this chapter, the additional static fields would typically include points of interest at locations other than at real exits (representing blocked exits, or misinformation) and/or would exclude points of interest at legitimate exits (representing exits that are unknown to an agent with that view of the world). For example, a static field incorporating belief in a non-existent (or blocked) exit would be calculated by supposing the presence of that exit and then carrying out the standard method for calculating the static field.

Having provided for multiple representations of the world, the agents must be modified to allow them to determine which representation is currently being consulted. Each agent is extended by adding an integer index into the set of static fields called the view selector, ψ . The view selector identifies the static field currently being consulted by the agent, which has no access to the information in other static fields.

The cell scoring function of the model (equation 5 of page 74), already modified in the last chapter (equation 9 of page 85) is altered to take account of the view selector:

$$c_{ij} = \exp(k_D D_{ij}) \exp(k_S S_{\psi_{ij}}) (1 - \eta_{ij}) \xi_{ij} \quad (11)$$

Notice that in equation 11 the term $S_{\psi_{ij}}$ replaces S_{ij} , reflecting the alteration in meaning of S from equation 10.

Equations 10 and 11 together represent the primary change made to the floor field model to support heterogeneous crowds based on knowledge of space. By providing a spatial representation at an appropriate scale — tightly integrated at the heart of the floor field model — we maintain the causal connection between these rules and the emergent crowd behaviours of the floor field model. This application of the Microscopic Human Factors methodology will allow us to characterise the effect of adding multiple representations to the floor field model, comparing it with the original model's single field results.

7.3.2 Physical discovery of spatial information

As the agent moves through the space it is necessary to provide for a mechanism that allows the agent to accumulate information, changing its beliefs about the world by updating its ψ attribute. The reduction provides for a spatial mechanism that identifies certain circumscribed areas as zones of detection for new spatial information.

In the ontology of the floor field model, the way to distribute information spatially is with floor fields, and the unit of resolution of all spatial information is the cell. This suggests a new field to manage discovery. Accordingly, we define a new floor field called the *discovery field*, I , in which the value on each cell (i, j) is an integer $0 \leq I_{ij} < |S|$ specifying the view selector that an encountering agent should adopt.² Encountering a cell (i, j) on which $I_{ij} \neq \psi$, implies a potential change to ψ to bring it in line with I_{ij} (see figure 17).

As described in the reduction, when encountering ψ' , a new index into S from the discovery field, agents must determine if this transition is allowed. This prevents agents from regressing when leaving information areas. (For example, in figure 17, if the agent steps to the west, into

²The symbol I (for information) is used as D already indicates the dynamic field.

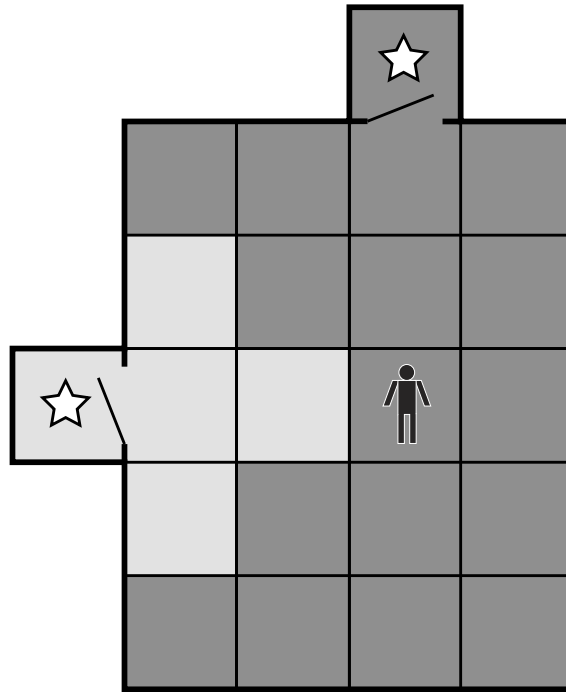


Figure 17: Discovery field, I , set up for a proximity effect in which approaching a previously unknown exit leads to discovery. In this example, the discovery field encodes two potential view selector values: 0 (*dark cells*) or 1 (*light cells*). These are potential agent ψ values, indices into the set S of static fields. Suppose that in this example, S_0 identifies distance only to the north exit while S_1 identifies distance to the closest exit. Suppose that the agent has $\psi = 0$, implying current consultation of S_0 and ignorance of the western exit; if the agent steps west, its ψ will be updated to 1, switching its view of the world to S_1 , and making it aware of the western exit. If the agent with $\psi = 0$ steps forward, it will remain unaware of the western exit, and may exit north (the only exit it knows about) never learning of the western exit.

the area with $I_{ij} = 1$ we want ψ to become 1. If the agent subsequently steps east, we do not want ψ to return to 0, as this would imply forgetting.) We have adopted a finite-state machine approach to this problem, specifying the allowable transitions $\psi \rightarrow \psi'$. For present purposes it is sufficient to enforce a very simple transition rule: ψ can only increase monotonically. Thus, in the present model:

$$\psi = \begin{cases} \psi' & \text{if } \psi < \psi', \\ \psi & \text{if } \psi \geq \psi' \end{cases} \quad (12)$$

In the present implementation we therefore ensure that the static fields appear in the set in order of increasing knowledge. In a two field scenario, for example, agents may initially believe there are two exits in the space, but upon attempting to use one of the exits discover it blocked. The former view represents a more naïve view and accordingly appears earlier in the set than the latter. The values of the discovery field, then, represent the minimum information available to an agent in that location, with ψ taking on the discovery field value of the local cell only if it represents an increase in knowledge.

7.3.3 Communication of spatial knowledge

Agents encountering one another can communicate spatial information. There is no prototype for this within the floor field model, so our direct communication model is a new development. Communication requires both a trigger condition, and a mechanism for information transfer. As discussed, the communication is triggered in the same manner as voluntary pushing: being blocked in movement.

Like pushing, communication is a directed activity from one agent toward another. Accordingly, a stigmergic mechanism involving an omni-directional field is a less desirable implementation compared to a more specific mechanism isolated to the two agents. Our implementation involves a direct message from a communicating agent to a blocking agent. Communication in this sense means sharing information in order to cause the blocking agent to be recruited to the communicating agent's goals; the communicating agent shares its view of the world with the blocking agent. As view of the world is represented by the ψ attribute, communication is

implemented by updating the ψ parameter of the blocking agent. Like the discovery mechanism, communication is for information gain; accordingly, equation 12 is again used to prevent communication-induced forgetting. The blocking agent makes use of the updated ψ parameter on the next time step.

7.3.4 Algorithm

Algorithm 3 is the evolution of the pseudo-code for the time step function first given in algorithm 1 for the floor field model, and already elaborated in algorithm 2 for the force simulation. In this algorithm, lines 9–10 of `STEP` provide the mechanism for discovering information (by consulting the discovery field, I , when moving to a new location) and for communicating to others (at the same time as force is delivered). In the algorithm, the `AGENT-AT` function returns the agent present on the desired cell and is considered primitive. The agent methods `SET-VIEW-SELECTOR` and `GET-VIEW-SELECTOR` are simple accessor methods whose implementation is not shown. Other procedures retain their implementation as given in previous chapters.

7.4 Results

The model scenarios described in chapter 6 were designed to evaluate the effects of forces on crowds, including crowds with high drive, resulting in extreme density and high forces that can be dangerous and cause injury. In this chapter our focus is on communication and the effect of heterogeneity in crowds. As such, the parameter values and scenarios presented are the ones that are the most revealing for this different purpose.

Three scenarios are presented in this section in order to show the effect of introducing discovery and communication in the model. First, we present a baseline scenario to show how the force-enabled model performs without any heterogeneity in movement goals. Second, we introduce a heterogeneous simulation. Finally, we enable communication.

Like the scenarios studied in chapter 6, scenarios in this chapter take place in a room of dimension 61×61 cells, surrounded by wall cells. Because we are less interested in bottlenecks and more interested in a crowd that moves enough to allow interactions between groups of

```

PROCEDURE STEP( $D, S, F, I, agents$ )
1  DECAY-AND-DIFFUSE( $D$ )
2  for each  $a$  in  $agents$ 
3      do  $a$ .CHECK-INJURY( $F$ )
4           $\psi \leftarrow a$ .GET-VIEW-SELECTOR()
5           $a$ .CELL-SELECT-WITH-FORCE( $D, S_\psi, F$ )
6  for each  $a$  in SHUFFLE( $agents$ )
7      do if  $a$ .ATTEMPT-MOVE() = FALSE
8          then  $a$ .PUSH()
9               $a$ .TELL( $\psi$ )
10         else  $a$ .DISCOVER( $I$ )
11  PROPAGATE( $F$ )

AGENT METHOD TELL( $\psi$ )
1   $s \leftarrow$  SELF
2   $(i, j) \leftarrow s$ .GET-DESIRED-CELL()
3   $r \leftarrow$  AGENT-AT( $i, j$ )
4   $r$ .HEAR( $\psi$ )

AGENT METHOD DISCOVER( $I$ )
1   $s \leftarrow$  SELF
2   $(i, j) \leftarrow s$ .CURRENT-LOCATION()
3   $s$ .HEAR( $I_{ij}$ )

AGENT METHOD HEAR( $\psi'$ )
1   $s \leftarrow$  SELF
2   $\psi \leftarrow a$ .GET-VIEW-SELECTOR()
3  if  $\psi' > \psi$ 
4      then  $s$ .SET-VIEW-SELECTOR( $\psi'$ )     $\triangleright$  New view selector not used until next step

```

Algorithm 3: Step function modified for discovery and communication

agents, we provide for many more exits in these scenarios. One wall of the room contains 10 exits, each of width 3 cells, spaced evenly along the wall (see figure 18).

We have continued to use the low and medium drive parameter values from chapter 6, but we have changed the high drive parameter values. As discussed in section 6.5.3, the parameter combination $(k_D, k_S) = (0, 10)$ produces extreme polarisation of agent cell preferences. As our interest here is not in extreme force situations, we have adopted $(k_D, k_S) = (0, 7)$ for the high drive case — parameter values that still produce relatively high drive, but largely eliminate the aisle effect. We shall have more to say on the subject of aisle effects in section 8.4.

As described in section 2.3, actions taken in crowd situations are best viewed in light of the information available to the individual at the time. Indeed, disaster investigations (e.g. Averill et al., 2005) often include questionnaire studies in which survivors are asked where and when they noticed danger cues and became aware of key facts; this information is of great interest in understanding behaviour. Accordingly, in the results reported here we have chosen to group agents based on ψ upon exit. This allows us to report how many agents in each information state were successful in exiting as well as how much distance these agents covered from their initial positions.

Although looking at the differences across groups is one way to determine the heterogeneity of the crowd, observation is key to understanding the behaviour of a heterogeneous crowd. By quantifying only exiting agents we naturally leave out agents who do not exit the space;




Figure 18: Physical environment for discovery and communication model scenarios. Ten exits shown (black cells in walls) with static field (gradient). Agents not shown.

analysis of the non-exiting agents is made more complex due to the fact that their failure to exit is sometimes produced by interactions between the groups. We will discuss the non-exiting agents by recourse to observation, reporting on them in the individual scenarios.

7.4.1 Homogeneous crowd

As a baseline, this first scenario contains only one static field that correctly identifies the locations of all the exit cells (see figure 18). All values of the discovery field indicate the use of this single static field, and all agents consult this static field from the outset of the model. Aside from the larger number of exits, this scenario, then, is identical to that of the previous chapter.

Figure 19 shows the same analysis as in figure 12 (page 104) but for our increased number of exits. By moving to a scenario with ample exits — rather than the bottleneck scenarios of chapter 6 — the pattern of injuries and exits has changed. As expected, the number of injuries has dramatically decreased and the exit rate is dramatically increased as compared with figure 12. Redefining the high drive parameters has not altered the superiority of the medium drive case on injuries and exits. (Welch's t -tests were used to compare the high and medium drive cases at each value of ϕ . In all cases the differences in means were found to be significant, $p < 0.001$.)

We found that with ample exits, the number of injuries — even at high drive when many agents became injured — did not have a noticeable effect on the exit rate in 350 iterations of the model. Because injuries do not have an interesting effect on movement, in this chapter we will generally report our results differently. We will only present exit results across values of ϕ where this variable interacted with the results of interest. In other cases, we use $\phi = 125$, the mid-point of the explored range, and we present results as shown in figure 20a. In this chapter, heterogeneous behaviour is of interest and so we also include a measure of steps taken by exiting agents (figure 20b, a count of the number of time-steps in which the exiting agent moved). This measure quantifies the directness of agent exits; by grouping agents by ψ at exit time, we can determine whether time spent in the model is a function of what an agent knows.






Figure 19: Agents exiting and injured with ample exits. Range of ϕ explored.



a. number of exiting agents



b. steps moved by exiting agents

Figure 20: Agent egress measures in a homogeneous crowd with plenty of exits ($\phi = 125$)

Let us consider figure 20b, the number of steps taken by agents at differing crowd densities. At high drive, the mean number of steps to exit is 30.36, very close to the 30.5 that would be expected for a complete optimal exit given a uniform distribution in a space of size 61×61 . The larger number of steps in the remaining two cases indicates a more circuitous route followed due to the influence of the larger k_D parameter, and consequent smaller influence of the static field (see sections 5.3 and 6.5.1 for more on this point). At medium drive, almost all the agents were able to exit, but many more steps were taken. At low drive, agents are not motivated to move toward exits, and so many fewer agents exit — primarily those initially located close to the doors; these agents, however, take very large numbers of steps before exiting (even if initially close to the exits).

We have established baseline results for our scenario with ample exit cells, and only a true static field. Agents are relatively unimpeded in exiting. Let us now turn to a question of discovery. What if agents believed in a second set of exits? How would this alter the dynamics of the crowd?

7.4.2 Heterogeneous crowd

This scenario takes place in the same physical space as the previous scenario. This scenario, however, contains two static fields. The first static field (figure 21), used by all agents at the outset of the model, reflects all of the real exits, but also supposes an identical set of exits




Figure 21: Static and discovery fields in the heterogeneous scenario. All agents start with the static field shown. The discovery field indicates in blue those cells that trigger a change to the second static field, which is identical to the one displayed in figure 18.

on the opposite wall (which are in fact blocked). The discovery field is configured so that agents approaching the blocked exits (i.e. coming within two cells of the wall with the blocked exits) switch to the second static field, which only reflects the good exits. For the purpose of characterising heterogeneity on its own, we here disable the communication behaviours in the model by disregarding line 9 of STEP in algorithm 3.

Analysis and understanding of the results depends on a description of the behaviour of the model agents. Observations of the model reveal that there are three categories of agents at the outset of the model:

1. agents unaware of the blocked exits (naïve agents with $\psi = 0$) located in the half of the room closest to the doors,
2. agents aware of the blocked exits (knowledgeable agents with $\psi = 1$) who are located within two cells the blocked exits, and
3. naïve agents located in the half of the room closest to the blocked exits.

These groups exhibit different behaviours and have different outcomes. The outcome for an agent is not only affected by its group membership, as there is also an interaction with agent drive. As we shall see, the freezing-by-heating paradox discussed in section 2.1.2 is key to understanding this interaction and is well demonstrated by the heterogeneous crowd in this scenario. We will first consider the model in the high and moderate drive cases. Then we will consider the low drive case and return to the question of freezing-by-heating.

High drive and moderate drive cases

Observation of the model (see figure 22) reveals the behaviour interactions between the three groups of agents, which differ only in degree between the high and moderate drive cases. Because of the influence of the static field in these cases the agents move rapidly toward the closest exit. Group 1 agents are free to move toward the true exits and exit readily, never becoming aware of the blocked exits. Group 2 agents (who are joined within the first few time-steps by a small number of group 3 agents close to the blocked exits) realise that they



Figure 22: Three groups of agents at discovery of blocked exits. Naïve agents exiting northwards (*group 1*). Note barrier formation as knowledgeable agents at south wall (*group 2, blue*) cannot pass southbound naïve agents (*group 3*) pushing toward blocked exits on south wall. $k_D = 0$, $k_S = 7$, $\phi = 125$, time step 11.

are at the wrong end of the room and begin to move toward the real exits. The remaining group 3 agents approach the blocked exits and encounter group 2 trying to move in the opposite direction.

In these cases, group 3 quickly forms a barrier for group 2 and vice versa. Because the exits are spaced evenly along the walls, the group 3 agents form into a band stretching from wall to wall; because of the large number of group 3 agents, the band is many agents deep, and (barring injuries) can physically push group 2 back toward the blocked exits. In the high drive case, the band forms rapidly, is high density and is impenetrable to group 2. In the moderate drive case, the band is slower to form, is less dense, and some group 2 agents are able to infiltrate themselves through the band.

Let us take up the question of injuries. Injuries (figure 23) followed the same pattern as in figure 12 in that higher drive to achieve goals results in more injuries. Injury rates are generally lower than in the scenario of chapter 6 because there are ample exits available along one wall, which accommodate many agents per time step. The porous wall reduces crowd density and hence movement conflicts at the exit. Note that the x -intercept of the high density case is at a lower value of ϕ in figure 23 than in figure 19; this occurs because enough agents are detained in the barrier formation to reduce the potential for injuries by further reducing density (and hence force) at the exits.




Figure 23: Agents injured in exiting from the blocked exit scenario without communication

Observation of the model shows remaining injuries are primarily concentrated where groups 2 and 3 meet. Injuries occur as group 2 agents (unable to make their way forward) exert force upon the tightly packed crowd of group 3 agents, and vice versa. As ϕ increases, injuries fall within the model as the agents are not numerous enough to bring a crushing force to bear (recall that group 1 agents exit quickly in all cases). Regardless of whether the situation is complicated by injuries or not, the stable barrier formation persists and a stalemate ensues in which neither group 2 nor group 3 exits the model.

This situation is reminiscent of the Lenin stadium disaster described in section 2.2.2 in which a smaller group colliding with a bigger one created a large number of injuries. It is also consistent with Fruin's observation that injuries within a crowd are not distributed at random, but follow patterns of force that accord with the context of the situation (Fruin, 1993).

This stable configuration — with or without injuries — explains the fact that exit rates are independent of ϕ in this scenario. Accordingly, as we now turn to the question of exits from the space and distance travelled we show only results for $\phi = 125$ for clarity.


The pattern of exits engendered by the model (figure 24a), together with our observations, confirm that group 2 agents ($\psi = 1$) are, for the most part, unable to exit in this scenario. In the high drive case only 1.3 knowledgeable agents on average manage to exit; they originate adjacent to the blocked exits and are able to move toward the real exits before the group 3 agents have fully blocked the route. In the moderate drive case the group 3 agents approach more slowly giving on average 32.8 agents a chance to escape the crush and move to the exit. The numbers of agents exiting with $\psi = 0$ in both the high and medium drive cases is the same because these are the group 1 agents (expected to number half of the 1116 initial agents) who have time for a complete exit in the high and moderate drive cases.

In terms of distance travelled (figure 24b), exiting knowledgeable agents with $\psi = 1$ required many more steps to exit than did naïve agents with $\psi = 0$. Note that agents exiting with $\psi = 1$ may have originated either in group 2 (knowledgeable at the outset) or in group 3 (naïve at the outset, but gaining knowledge through discovery).


Low drive to exit

Model behaviour when drive to exit is low is quite different from the previous cases. In this case, the low influence of k_S relative to k_D means that agents display only a weak movement bias toward areas believed to be exits. The classification of agents into three groups is still meaningful in terms of explaining their behaviour, but the dramatic conflict between groups 2 and 3 does not occur. This is because the group 3 agents do not move toward the blocked exits in large number, leaving large spaces that allow group 2 agents the freedom to infiltrate toward the real exits. Indeed, the space created by moving group 2 agents allows agents from group 3 to move into the blocked exit area, discover the blockage, and join the group 2 agents in moving toward the open exits.

Referring to figure 24, however, it is clear that avoidance of a stalemate between groups 2 and 3 does not translate into exits by knowledgeable agents with $\psi = 1$. This is because the agents, even though they are knowledgeable about blocked exits are simply not driven to move



a. number exiting. (Mean over 50 replications)



b. steps taken by agents exiting. Note that this measure is undefined if no agent exits in a particular replication. Here we calculate the mean across the n of 50 replications in which agents exited)

Figure 24: Exiting agents in the blocked exit scenario without communication. $\phi = 125$

toward the real exits. The exit and step pattern for low drive agents in figure 24 is essentially the same as that shown in the baseline case of figure 20.

Section 7.4.1 determined that in a scenario with the same number of real exits as this one egress was not impeded and minimal injuries occurred. In the present scenario, simply by introducing belief in a set of blocked exits, the heterogeneity of the crowd resulted in a much lower exit rate and a much higher injury rate. In this case the conflict occurred due to agents working at cross-purposes. Agents that had discovered information moved one way, while agents who had yet to discover the information moved in the other. In this scenario, the communication potential of the agents was suppressed, and so physical discovery of the information was a necessity. Let us now determine the value of inter-agent communication by activating the communication simulation within the model.

7.4.3 Communication in a heterogeneous crowd

In this scenario we take up exactly the same situation as in section 7.4.2, except that we now enable the communication functionality of the model (line 9 of *SIM-STEP* in algorithm 3). In cases where they are blocked in movement, agents will now communicate their view selector to the blocking agent.

In this scenario, the same three groups arise at the outset of the model as occurred in the last scenario. As the simulation unfolds, however, the barrier formation where naïve group 3 agents and knowledgeable group 2 agents jam one another in place does not occur. This is because when a naïve agent moves to block a knowledgeable one, the knowledgeable agent communicates its view of the world, updating the ψ value of the naïve one; this has the effect of explaining the blocked exit condition to the naïve agent, turning it into a knowledgeable one. Both agents move off toward the real exits, informing other conflicting agents on the way (we shall discuss the issue of trust below).

The injury and exit rates in this model demonstrate the value of communication. Injury rates for this scenario (figure 25) are much improved from the non-communicating case (figure 23





Figure 25: Injuries with communication enabled


in section 7.4.2), to the point of being comparable (if not identical) to the results of the homogeneous scenario (figure 19 in section 7.4.1) in which all agents were knowledgeable.

When agent drive is high, injuries are actually improved slightly from the baseline case reported in section 7.4.1; this occurs because the larger number of steps for agents in group 3 has the effect of delaying their arrival at the exits, consequently reducing density and pressure in that area. This effect, however, becomes minor as agents become more resilient: we note that the x -intercept at high density has been restored to approximately the same value as in figure 19; now that the barrier formation has broken down, the pattern reverts to one in which all agents are knowledgeable and hence density, force and injuries return to the same as in the homogeneous case.

Figure 26 shows exit results broken down by ψ value. In interpreting this graph we must bear in mind that communication causes transitions between the ψ groups at exit time if there is a bottleneck. This is a consequence of our abstract trigger mechanism that presumes agents communicate whenever blocked in movement. Knowledgeable agents arriving at a bottleneck



a. number exiting.



b. steps taken by agents exiting.

Figure 26: Exiting agents in the blocked exit scenario with communication. $\phi = 125$

will communicate with the rearmost agents at the exit. These agents in turn will communicate with those in front. The word travels quickly in a dense crowd, and soon all agents exiting from the bottleneck are counted as knowledgeable agents. Similar challenges relating to determining who knew what when also occur in investigations of real crowd incidents, and explain the focus of investigators on *when* people became aware of various aspects of the situation. This level of data collection is beyond the scope of what is reported here, leading to watering down of differences between groups.

This comment is important in understanding the appearance of even exit rates between ψ values in figure 26a. Some exiting group 1 agents actually exit with $\psi = 1$. Although the majority of agents left after 350 time-steps originated in the far half of the space, they recruit enough agents to skew the figures. Despite this, figure 26b shows that heterogeneity is preserved in the crowd. Even with some watering down of differences between groups, agents with $\psi = 1$ still take approximately three times as long to exit on average in these cases. Once again, agent drive accounts for the reduction in exit rate for moderate and low drive cases.

7.5 Discussion

Freezing-by-heating. Let us now return to the question of freezing-by-heating. Section 2.1.2 discussed this paradox in terms of two groups of agents wishing to pass through one another in opposite directions, and we alluded to its demonstration in the force simulation in section 6.8. In section 2.2.2 we reviewed the Lenin Stadium disaster that exemplified freezing-by-heating. When we introduced differing agent goals at high drive through discovery in section 7.4.2 we replicated this effect (the interaction between groups 2 and 3 providing a direct simulation): At low agent drive, crowd densities and speeds remain low, allowing for an accommodation as groups 2 and 3 infiltrate themselves through the gaps in the opposing group. When drive becomes moderate, the density increases and there are many fewer empty spaces in the band formed by group 3. This leads to interference between the groups and limited potential for group 2 agents to access the real exits. In the high drive case, the group 3 agents essentially form an impenetrable wall, and the group 2 agents are pinned behind it. Thus, as the energy of the

system increases from the low drive to the high drive case, the resultant pattern of interaction between the groups becomes more rigid, culminating in a frozen stand-off when energies are highest. By introducing heterogeneity into a crowd movement model, we have been able to demonstrate this effect which is not seen if all agents behave identically in their move toward the exits.

We note that jams in counter-flow have been directly studied in the floor field model (Burstedde, Kirchner, Klauck, Schadschneider, & Zittartz, 2002), these jams were quite simple. They were orchestrated by directing equal-sized fixed groups of agents at one another and do not involve force or any kind of 'heating.' Our study of freezing-by-heating involves emergence of groups due to dynamic discovery of spatial information, emergence of jamming, force, injuries and communication. This is a more realistic scenario and a better demonstration of the principle.

We note that the freezing-by-heating paradox adds an additional argument to the urgency of calm and measured movement for the sake of safety in crowded conditions. It has been noted in fire disasters, however, that slow response times waste valuable time early in an emergency (Sime, 1995) which may leave insufficient time for an orderly and measured egress. Our results suggest that competitive increases in drive can result in a frozen crowd.

Situatedness and knowledge representation. Let us now consider the heterogeneity of the crowd. In a sense, the original floor field model can be seen as supporting a heterogeneous crowd because the static field can break a crowd into sub-crowds that move toward differing targets. This occurs because the floor field model naturally provides for a situated simulation. In keeping with the concept of situated cognition in cognitive science (Smith, 1999), agents directly consult their immediate surroundings (the local neighbourhood in the static field) regularly, making local movement decisions based on current conditions (including the movement of other local agents); this can be contrasted with a traditional artificial intelligence approach to planning that pre-computes complete and optimal paths in an abstract planning environment, with the actual path traversal delegated to a separate system in a distinct process. The heterogeneity present in our model goes beyond the sub-crowd-based heterogeneity of the floor field model.

This is because the sub-crowds in the floor field model are themselves homogeneous. We have introduced a mechanism for agents to be part of the same crowd, in the same physical area, and yet not be focused on the same objectives.

The differences in exit performance by agents with different information states (represented within the model by different values of ψ) does quantify a pattern of heterogeneous behaviour within the crowd. We note, however, that differences *within* information states can be as important as differences *between* such states. For example, agents with $\psi = 0$ are divided into groups 1 and 3 because of situated information processing despite their identical mental content. This difference between agents that have the same mental content demonstrates how the MHF approach to integrating cognitive processing preserves the situatedness of the floor field model. By integrating knowledge representation at the level of movement rules within the individual, our method produces agents that are highly situated.

7.6 Model validity

Implementation. At the implementation level we compare the behaviour of the model to the intended abstract crowd behaviour specified in the reduction. The differences in performance for agents with distinct ψ , as well as observations of the differing behaviour of these agents, indicates that the model has implemented the mental content requirements of the reduction. Discovery by agents occurs upon entry to specified areas within the modelled floor space, and the multiple static fields and equation 12 provide for the accumulation of knowledge. Coupling directed communication with the pushing mechanism, a straightforward change given the implementation of chapter 6, has been verified by observation.

The other task at the implementation level is to determine if the implementation formalism is appropriate for the reduction. In general we have found the floor field model, as augmented by our force simulation, to be a natural implementation medium for the reduction. As discussed in section 7.3.2, the state transition rule to increase ψ strictly monotonically to model knowledge accumulation has been adequate for the investigations here; a more elaborated finite-state implementation would provide for more flexibility where required.

Reduction. At the level of the reduction we ask whether the rules provided were sufficient to implement the effects we described in the specification. The specification was derived from Sime's description of a scenario as an information system distributed through a physical environment. Experiments with the model demonstrate that, indeed, the reduced rules transform the space of the floor field model into an information system capable of radically altering agent behaviour upon discovery of new information. We also noted in the specification that in a heterogeneous crowd, physical proximity is not a good predictor of similar behaviour; this has been amply demonstrated by the model, for example, when group 2 and group 3 agents interact.

Introduction of communication into the model has added additional richness to the information system, allowing agents to become aware of their environment without physically visiting locations of interest. As we intimated in section 7.2, communication is a highly complicated behaviour that is deployed to many different ends and is triggered in many different ways. A more complete specification of communication itself, and subsequent modelling to investigate its mechanisms is accordingly left to future work.

Specification. Ultimately, validation at the level of the specification asks what the model can tell us about crowd behaviour. This demonstration was different in goal from that of the previous chapter. Here, we decided to directly provide for a heterogeneous crowd to investigate the interactions between the groups created within the crowd. As stated in the specification, it was not our intention to model all possible causes of heterogeneous crowds, or all factors in their maintenance. We found, however, that modelling discovery and communication did lead to a heterogeneous crowds in which force dramatically affected movement in the model; heterogeneity totally changed the ability of agents to move freely within the space, results that accord with reductions in capacity due to crossing flows (Still, 2000), and reproduced the type of crowd disaster that occurred at Lenin stadium (section 2.2.2).

Introducing communication into the model increases the chances that adjacent (conflicting) agents will begin to act in the same way (by sharing world views). Our number of steps measure, however, demonstrates continued heterogeneity, as the time-line of knowledge acquisition for

each agent influences their pattern of movement over the life of the simulation — and ultimately time to exit in the scenarios studied.

Introduction of discovery and communication has, then, improved the model with respect to Sime's concerns. By making agents more relevant to real people through the use of the MHF methodology, we have shown agents that are no longer interchangeable, as their personal history within the model is now relevant to understanding their behaviour and performance. There is ample room for further work in this regard, as egress researchers have begun to consider the role of occupant experience with physical structures in understanding behaviour. The 'fire and ICE' concept (Gwynne & Boswell, in press), for example, argues that learning by occupants, by exposure to information during ingress and circulation, is crucial to understanding what they do in emergency egress. This observation suggests that egress models must become more general, simulating occupants in everyday use of a space as well as in egress, and allocating to exiting agents the knowledge acquired during the exploration/occupancy phase (Gwynne & Kuligowski, in press). An abstract investigation of this approach could use the model described here as a starting point, and we leave this interesting possibility for future work.

7.7 Summary

We have extended the floor field model further from the force model of chapter 6, adding a model of discovery, communication and representation of knowledge. Basing our specification on Sime's observation that crowd situations are information systems through which people move, we have allowed agents within the model to obtain spatial information from the environment as they circulate through the space.

Although we did not set out to model all possible factors that underlie crowd heterogeneity, we conclude that our premises do generate a heterogeneous crowd, although not necessarily along the lines of information state alone; the geometry of the space is also relevant, as agents are situated and take action based on local information, despite having more general knowledge. Obtaining new information leads to new movement goals, and, depending on agent drive, conflicts between agents of varying severity. This led, at high agent drive, to a barrier formation,

a reproduction of freezing-by-heating behaviours that echo the Lenin stadium disaster described in section 2.2.2. The simple impetus for heterogeneity in this model leads to the possibility of investigating other factors that lead to heterogeneity, and we leave this process to future work.

Although it may seem evident, it is worth noting that the ability to represent heterogeneous crowds and differing mental content allows for modelling more complex and realistic scenarios. Although we can fix mental content at the outset of a simulation, this allows for studying only small scenarios in which agents are interchangeable and not particularly realistic. As we have seen in our blocked exit simulation, and as seen at Lenin stadium, many situations unfold in several phases and without being able to represent changing goals we are unable to represent this within the model. If a model cannot represent these non-trivial — but interesting — scenarios, it misses out on discovering the interesting interactions between groups in the simulation. We offer our additions of discovery and communication, and their effect on evolving mental content, as a contribution in this regard beyond their role as simply demonstrations of the MHF methodology.

We found that communication in crowds can be a powerful mediator for crowd forces when it brings about a change in behaviour that is more complementary between groups. We found that it eliminated the barrier formation in the model, resulting in dramatic increases in exits and decreases in injuries. As noted above, communication is much more complicated in the world than we have represented here. Altruism and leadership (role) would be two examples of communication that are left to future work. Several parameters to communication are also possible, for example how likely a person is to engage in communication when blocked, or to believe a communication when received. We leave these parameters to future work as well.

One question is raised by the positive effect of spatial communication in dissolving barrier formations like the one that occurred at Lenin Stadium. Would improved knowledge (e.g. through spatial communication) of the spatial aspects of that situation have helped crowd members avoid a disaster? We suspect that the Lenin Stadium disaster did not include any spatial disagreement among crowd members. After all, the exiting people knew the direction to the exits, and the re-entering people knew the direction back to the field. Since spatial communication dissolves a barrier formation only by resolving disagreements about spatial facts, it does not help situations

that do not have any disagreements about spatial facts. It is simply the wrong tool for the job in such cases.

In chapter 8, we shall consider a different type of communication: front-to-back communication. This type of communication is not oriented at sharing spatial information. Rather, it focuses directly on reducing the crowd forces that cause injuries. We shall return to the question of the Lenin Stadium disaster when we consider the effects of this different type of communication on barrier formations in section 8.6.

Chapter 8

Front-to-back communication

This chapter forms the third demonstration of the use of the Microscopic Human Factors methodology. This demonstration differs from the demonstrations of chapters 6 and 7; while those chapters studied well-known (or, at least, accepted) mechanisms for crowd behaviours, this chapter uses the methodology to more prospective ends: exploring a mechanism that has not been previously described in detail.

We have seen in section 2.1.2 and chapter 6 that crowd disasters frequently involve high pressures that cause crushing injuries and death by compressive asphyxia. One might ask why people continue to apply pressure in these situations, given its disastrous effects. It has been suggested that failures in front-to-back communication are common in crowd disasters (Fruin, 1993; Pauls, 1984). These failures occur when those applying force at the rear of a crowd do not know of important conditions (e.g. blocked exits, fallen people) at the front of a crowd. As forces become large, those at the rear also do not know that their actions are injuring those at the front. This may explain continued pushing behaviours in these disasters.

A good example of a failure of front-to-back communication occurred at the Who concert disaster at Riverfront Coliseum (see section 2.2.3). Johnson's analysis of the event suggests that people pushing at the rear of the crowd were unaware of problems at the front (Johnson, 1987b). Although the media described the event as a stampede driven by mob psychology, Johnson found the opposite, that helping behaviours were widespread, even between strangers.

Although front-to-back communication failures have long been suspected as a significant factor in the unfolding of a crowd disaster, we have not found in the literature a systematic investigation of front-to-back communication in crowds, neither its mechanisms, benefits or model studies. This means that although we hypothesise that increased front-to-back communication would be helpful, we do not know the circumstances under which front-to-back communication is possible or what benefit could be practically expected from employing it. We do not know whether situation-specific factors affect front-to-back communication. This demonstration of the methodology represents the first steps at characterising this crowd-level behaviour, investigating some of these issues within the context of a microscopic crowd model.

The organisation of the chapter proceeds as follows. We first develop a specification for front-to-back communication itself, as it has not previously been formally described, and provide a reduction. We then turn to the details of our implementation of this reduction. We use the resulting simulation model to explore front-to-back communication in a simple laboratory-like scenario. Wishing to consider how a more realistic scenario affects front-to-back communication, we reconstruct the Who concert disaster in the model, making plausible observations concerning mechanisms of this disaster, and looking at those mechanisms when front-to-back communication is added. We return to the question of communication at Lenin stadium before closing with some preliminary conclusions on the nature of front-to-back communication in crowds, and proposals for further research work.

8.1 Specification

We are not aware of a study of front-to-back communication in the literature — either a formal review of case studies or a theoretical discussion about how front-to-back communication works within a crowd. Consequently, as the behaviour is not well understood, we here provide a hypothesis about how front-to-back communication *could* work in a crowd.

As noted in chapter 7, the person in the crowd is part of an information system. There is no reason, however, to consider this system to be composed only of the inanimate surroundings, or only to be concerned with physical structure. Information concerning the emergent crowd

conditions can also be communicated, and this is what is meant by front-to-back communication. Specifically, we are here concerned with force and injuries. In our view, there are four processes at work in effective front-to-back communication in a real crowd. These are: initiation (the creation of the front-to-back communication signal), propagation (movement of the message from front to back), direct action (helping behaviours) and decay (resumption of normal behaviour). We shall address each of these processes in turn:

Initiation. People likely employ several cues in a decision to initiate communication. These could include crowd density, movement, experienced force, noise level, visual observations, and other factors. By analogy to the bystander effect (helping behaviour in an emergency is inversely proportional to the number of bystanders, c.f. Latané & Nida, 1981) that is reduced when situations appear more dangerous (Fischer, Greitemeyer, Pollozek, & Frey, 2006), we assume that social pressure to appear poised in an emergency will decline in importance as perception of physical danger increases. Accordingly, in specifying front-to-back communication we highlight the importance of perception of safety in the decision to communicate.

Propagation. When people get the message that those in front are experiencing negative effects, we hypothesise that people can do things to reduce the negative effects. One simple action that seems generally possible even in a dense crowd is repeating the message toward the rear. As in chapter 7, we acknowledge reports that yelling a distance of 10 feet was impossible due to the noise at Riverfront Coliseum (Johnson, 1987b). Thus, we take front-to-back communication to be a distributed process of communication that involves direct personal interactions. As in other confusing or noisy environments, successful communication requires two things: first, obtaining the attention of another individual, and second, directly communicating a simple message to them. Propagation of front-to-back communication through the crowd depends on consecutive repetition of the information over time. We suspect that the key to successful propagation is a simple message. We feel a simple message, like “tell the people at the back to stop pushing” or “people are getting hurt in front” is most likely to be repeated over a distance without being

distorted. (We shall consider the question of successful propagation in section 8.8; see also the discussion of section 9.1.)

People in large crowds do not have an overall view of an unfolding crowd situation and may not know their exact position. They are restricted to diffuse sources of information like smells or noise levels, or purely local information like level of experienced force. They cannot be expected to know *why* others push or move in particular directions.¹ Interpretation of cues may be especially difficult in a complex situation with distributed focal points, heterogeneous crowd goals and multiple points of attraction. As previously discussed, the behaviour of individuals is best understood with respect to the information available in the local environment at the current time. Accordingly, we presume that, when we speak of front-to-back communication, 'back' is not a global concept, and can only be determined locally. Two people, even in close proximity, may conclude that the back of the crowd lies in different directions.

Helping behaviours. Fruin tells us that in a densely packed crowd, individual control can be lost (1993), as seen at Riverfront Coliseum where people were powerless to avoid the fallen (Johnson, 1987b). Thus, many people within a crowd have limited ability to take action beyond the propagation just described. Some people, particularly those not experiencing high forces, have the capability to act more directly by taking action to mitigate the negative effects experienced by crowd members further forward. Johnson reported a willingness to help (1987b); helping actions could range from complex (e.g. contacting security personnel) to simple (relaxing goals, increasing patience and decreasing applied forces).

Decay. We hypothesise that provision of helping behaviours is a temporary process that is self-limiting. People in crowds are independent, and in the case of moving crowds, have an objective that they are trying to achieve. As they do not have access to global information, a decision to help does not imply knowledge of how long to help for. In the absence of other

¹Knowledge of others' intentions is an interesting problem in agent-based modelling, and has been studied generally in the context of economic games like the prisoner's dilemma (e.g. Axelrod, 1984). This is related to the philosophical question of 'theory of mind,' in which we try to determine and understand what others believe (see e.g. Gopnik, 1999).

information, we hypothesise that people will return to their normal behaviour. Should future communication arise, helping behaviours may arise again.

By specifying our front-to-back communication mechanism in the absence of data from real crowds, we leave open the possibility that these processes may not completely capture the behaviour of front-to-back communication in real crowds. However, our goal is a qualitative investigation (a quantitative model would require model parameters for which experimental measures have yet to be forthcoming). We believe that this specification is plausible for the purpose of evaluating the potential benefits of front-to-back communication, as well as the viability of person-to-person directed communication within a crowd.

8.2 Reduction

The goal of the reduction is to take our hypotheses about front-to-back communication and to create an abstract description of them: essential situated rules specified from the point of view of an individual within the crowd. Here, we provide a description of each process identified in the specification:

Initiation. A person consciously initiates front-to-back communication, based on stimuli that are cognitively available in the local environment. As discussed in the specification there are several types of cues that can potentially be considered, but we put a priority on personal safety cues. The primary threat to personal safety in crowds is extreme forces (a belief that has motivated the modelling work of chapter 6). Thus, in our model, front-to-back communication is initiated in response to a perceived threat to safety from an experienced strong force (e.g. resulting from external restriction on an agent's desired actions). In other words, we suppose that people use the force they are experiencing in the crowd as a source of information, and use this information as the basis to initiate front-to-back communication. We feel that as other types of cues (e.g. visual or auditory ones) would hasten rather than delay communication, leaving their investigation to future work is a conservative decision in evaluating the effectiveness of front-to-back communication.

Locating the rear. As discussed in the specification, only local, or general non-specific information (like sound and smell) can be used in determining direction for communication. In initiation, as in retransmission (below), we use experienced force as a source of information; the direction opposite to the experienced force is taken to be the back of the crowd. There are three reasons for this decision. First, front-to-back communication is about reducing force upon agents at the front; using force as a predictor of the direction rear seems like a plausible choice for this type of message. Second, as with initiation, using a simple local mechanism like force is conservative. Third, as we are using force for initiation of communication, in an abstract model using the same source of information for propagation is parsimonious. Note that as force has an emergent evolution with time, people's concept of the direction toward the back will also change with time.

Retransmission by dyads. Just as in chapter 7, we take this form of communication to be an interaction between two individuals. In order for front-to-back communication to be effective, this process must be repeated from one person to another in the crowd. Like force, information takes time to propagate. We cannot assume that those at the back will immediately know about communication initiated at the front. In an abstract model of front-to-back communication the exact content of the message is less important. Accordingly, we speak of a simple *signal* that is propagated. This could consist of a simple spoken message, or even a physical signal assuming it were to be generally understood.

We assume that force must reach a certain level before retransmission is possible. This is because in the absence of a perceptible force gradient agents cannot determine the rear as described above. We hypothesise that frequent small jostles are prevalent in crowds, this is a noisy signal that we do not want to propagate, or use in behavioural decisions (like determination of the 'back').

Helping behaviours. The specification acknowledges that agents who are in low-force areas can potentially execute helping behaviours. Complex behaviours (e.g. reporting a problem to authorities) are outside the scope of investigation using a microscopic model. Simple helping

behaviours are nevertheless possible. As we have seen in chapters 6 and 7, by reducing drive to reach goals crowds can become safer. In addition, we assume that people who are capable of voluntarily pushing can cease to do so, although their involuntary leaning forces cannot be controlled. Accordingly, helping behaviours in the model include cessation of voluntary pushing and drive reduction.

Decay. The time-based mechanism of the specification is directly adopted by the reduction. After a delay agents will discontinue helping behaviours and will return to their normal behaviour.

8.3 Implementation

Building on the model of chapter 6, we have implemented the rules of the reduction in the context of the floor field model with our force additions. The four key processes of initiation, action, retransmission and decay are implemented as follows in our front-to-back communication simulation. An algorithm for these processes is given in section 8.3.1.

Initiation. Just as in chapter 7 where communication of spatial information was linked to an existing agent behaviour (in that case, to frustration in movement) we have linked the initiation of front-to-back communication to another existing agent threshold. The reduction specifies that force and a threat to safety are the triggers for front-to-back communication. We believe that the force at which agents lose control of their own movement is a salient and worrying level of force to experience. Thus, we couple initiation of front-to-back communication to the χ force threshold at which agents are forced to select the cell in the direction of incoming force (i.e. when an agent perceives a loss of decision making capability).

Similar to the communication mechanism of chapter 6, a direct agent-to-agent communication is used. Agents deliver a simple signal toward the back. In accordance with the discussion of local information in the reduction, ‘back’ is defined locally at the time of communication as the direction opposite \vec{f}_{ij} , the vector sum of force bosons present on the agent’s cell (equation 7 on page 84), and is quantised using equation 8. Agents who initiate front-to-back communication

and continue to experience force above the threshold may initiate communication again on the next time step.

Action. An agent initiating or accepting front-to-back communication will take action in two ways. First, the cell selection function (originally specified as equation 5, modified as equations 9 and 11) is further adapted by introducing a new factor multiplying k_S :

$$c_{ij} = \exp(k_D D_{ij}) \exp(m_S k_S S_{\psi_{ij}}) (1 - \eta_{ij}) \xi_{ij} \quad (13)$$

When the signal is received, the new m_S factor changes from 1 to 0. This dynamically eliminates the receiving agent's desire to move toward exits by eliminating consideration of the static field. As we have seen in chapters 6 and 7, this reduction in agent drive causes a slower move toward the exit, reduces crowd density and can reduce injuries. The second action taken is that agents refrain from voluntary pushing when blocked in movement. Involuntary pushing continues as normal, as does normal propagation of existing force. (The communication of spatial information of chapter 7, if applicable, also continues as normal.)

Retransmission. An agent that receives one or more front-to-back communication signals in time step t will consequently retransmit up to once (one signal to one other agent) during time step $t + 1$ only. Future retransmissions require further signal reception; as initiating agents may initiate consecutively, this further signal reception is not unlikely. Retransmission only occurs if the scalar force present on the retransmitting agent's cell exceeds its pushing force, ρ (taken as the perceptual detection threshold for force), and if a non-injured agent exists in the selected direction. In accordance with the local information premise, a retransmitting agent reverses the force vector on its own cell to determine which way appears 'back'. Communication then occurs identically to initiation.

Decay. Agents acting upon front-to-back communication signals have a probability, γ , evaluated each time step, of returning to normal behaviour. Normal behaviour involves a resumption of voluntary pushing, as well as normal cell selection through $m_S = 1$.

8.3.1 Algorithm for front-to-back communication

The time step function, already elaborated in previous chapters, is modified again for front-to-back communication, and given by algorithm 4. Following the pattern already established, the procedures `GET-SIGNAL-STATE` and `SET-SIGNAL-STATE` are simple accessors whose implementation is not shown.

The requirement that retransmission should be delayed to the time step after receiving a signal requires the introduction of a new agent property *signal-state* that operates like a finite-state machine, taking a value from the following set: $\{\text{NIL}, \text{DELAY}, \text{ALLOW}, \text{ALLOW-AGAIN}, \text{DONE}\}$. All values except `NIL` indicate the agent is currently honouring the signal and will refrain from voluntary pushing (line 9) until decay of this state (line 14). The normal pattern is that agents receive a signal and enter the `DELAY` state. Upon the next time step `DELAY` becomes `ALLOW`, the signal is re-transmitted, and `ALLOW` becomes `DONE`. The `ALLOW-AGAIN` state indicates that a signal was received during the `ALLOW` state, implying that retransmission should be done in the current time step (like for `ALLOW`) and also in the next time step (like for `DELAY`).

For clarity and brevity, additional notation has been introduced in calls to the agent method `RESET-SIGNAL-STATE` (primitive, implementation not shown). This method accepts a list of implications. If any antecedent is equal to the agent's current *signal-state*, then the signal state is assigned its consequent. If no antecedent matches then no change is made to the *signal-state*. Thus, the statement `RESET-SIGNAL-STATE($a \rightarrow b, c \rightarrow d$)` is equivalent to:

```

1  if SELF.GET-SIGNAL-STATE() = a
2    then SELF.SET-SIGNAL-STATE(b)
3  elseif SELF.GET-SIGNAL-STATE() = c
4    then SELF.SET-SIGNAL-STATE(d)

```

8.4 Laboratory scenario

We have evaluated the proposed front-to-back communication model using three different scenarios: a simple laboratory-type scenario, a realistic scenario (Riverfront Coliseum) and a scenario reminiscent of the Lenin stadium disaster. The laboratory scenario adopts the same scenario and parameters used in chapter 6 in order to determine whether front-to-back communication alters

```

PROCEDURE STEP( $D, S, F, I, agents$ )
1  DECAY-AND-DIFFUSE( $D$ )
2  PREPARE-AND-INITIATE-SIGNALS( $agents$ )
3  for each  $a$  in  $agents$ 
4      do  $a$ .CHECK-INJURY( $F$ )
5           $\psi \leftarrow a$ .GET-VIEW-SELECTOR()
6           $a$ .CELL-SELECT-WITH-FORCE( $D, S_\psi, F$ )
7  for each  $a$  in SHUFFLE( $agents$ )
8      do if  $a$ .ATTEMPT-MOVE() = FALSE
9          then if  $a$ .GET-SIGNAL-STATE() = NIL or  $\dot{f}_{ij} > \chi$ 
10             then  $a$ .PUSH()  $\triangleright$  voluntary pushing bypassed if honouring signal
11                  $a$ .TELL( $\psi$ )
12             else  $a$ .DISCOVER( $I$ )
13             if RANDOM()  $< \gamma$   $\triangleright$  decay the action of front-to-back communication
14                 then  $a$ .SET-SIGNAL-STATE(NIL)
15  for each  $a$  in  $agents$ 
16      do if  $a$ .GET-SIGNAL-STATE() in {ALLOW, ALLOW-AGAIN}
17          then  $a$ .TRANSMIT()  $\triangleright$  retransmission of signal
18  PROPAGATE( $F$ )

PREPARE-AND-INITIATE-SIGNALS( $agents$ )
1  for each  $a$  in  $agents$   $\triangleright$  unlock retransmissions from last step
2      do  $a$ .RESET-SIGNAL-STATE(DELAY  $\rightarrow$  ALLOW)
3  for each  $a$  in  $agents$ 
4      do  $(i, j) \leftarrow a$ .GET-CURRENT-CELL
5          if  $\dot{f}_{ij} > \chi$ 
6              then  $a$ .TRANSMIT()  $\triangleright$  Initiate front-to-back communication
7                   $a$ .SET-SIGNAL-STATE(DONE)

AGENT METHOD TRANSMIT()
1   $(i, j) \leftarrow s$ .GET-CURRENT-CELL
2  if  $\dot{f}_{ij} > \rho$   $\triangleright$  Agent pushing force is minimum for determining direction
3      then  $(i', j') \leftarrow$  CELL-IN-DIRECTION( $\hat{f}_{ij} + 180, i, j$ )
4          AGENT-AT( $i', j'$ ).HEAR-SIGNAL()
5   $s$ .RESET-SIGNAL-STATE(ALLOW  $\rightarrow$  DONE, ALLOW-AGAIN  $\rightarrow$  ALLOW)

AGENT METHOD HEAR-SIGNAL()
     $s$ .RESET-SIGNAL-STATE(DONE  $\rightarrow$  DELAY, NIL  $\rightarrow$  DELAY, ALLOW  $\rightarrow$  ALLOW-AGAIN)

```

Algorithm 4: Step function modified for front-to-back communication

the patterns of injuries or exit time. We set the probability of decaying helping behaviours to $\gamma = 0.1$, giving an expected duration of 10 time-steps. This is a moderate value for γ , and we leave a complete exploration of this parameter's value to future work.

8.4.1 Effect on injuries and exits

Figure 27 compares the results of chapter 6 to the same scenario with front-to-back communication enabled. We find that introduction of front-to-back communication into a laboratory-type scenario reduces injuries. This occurs regardless of agent drive. An important question that arises from these results is how exit rates are affected by the reduction in injury rates. Figure 27 indicates that front-to-back communication does affect the exit rate.

At low drive, agents are primarily guided by the movement of others. In this case the drive to exit and crowd density remain low, injuries are not a significant factor, and front-to-back communication further slows what is already a non-urgent exit from the space.

When agent drive is moderate, crowd density is also moderate. Introducing front-to-back communication improves the exit rate when ϕ is low; as agents are easily crushed in this case, front-to-back communication reduces forces, postponing injuries, and allowing more agents through the exit before it becomes clogged with injuries. When agents are more robust, numbers of injuries drop, revealing that — in the absence of injuries — the tendency of front-to-back communication to produce a more patient crowd results in lower exit rates.

Unlike the low and moderate drive cases, the high drive case shows an alteration of the multi-phase interaction with ϕ (discussed in section 6.7.1). The new relationship has a functional form more similar to the medium drive case, with many more agents exiting. Observations of the model suggest that adding front-to-back communication to the high drive case reduces the strong forces that usually accompany fast movement toward the exit and consequent high crowd densities. We observed that front-to-back communication both reduced injuries and also resulted in a faster exit. Let us examine these results further.




Figure 27: Effect of front-to-back communication in a laboratory scenario. Upper results for injury, lower results for exits. Lines identified by \times are copied from figure 12. Lines identified by $+$ indicate results with front-to-back communication.

8.4.2 Effect on the aisle

This third case, in which the exit rate is facilitated, is quite interesting in the context of the results of section 6.5.2 concerning the aisle region. In that analysis we found that high pressures in the crowd led to a formation called the aisle — a region centred on the door extending from the front to the back of the crowd in which agents are able to exit even under high pressure (figure 28a). By contrast, agents adjacent to the aisle are not able to choose a lateral step into the aisle because the forward pressure upon them exceeds χ and they are pinned in place. A stable configuration of this sort greatly decreases the exit rate. Recall that when injuries were introduced into the model an interaction emerged between ϕ and the exit rate in which moderate ϕ values (circa 95 in figure 27) provoked moderate numbers of injuries that acted as force-breaks within the crowd. These force breaks shortened the life of the aisle pattern by allowing individual action closer to the front.

When we introduce front-to-back communication into the model the aisle pattern disappears entirely. The results of figure 27, along with observations we have made of the model, suggest that front-to-back communication is an alternative to force breaks in facilitating exits during high-pressure egress. When front-to-back communication occurs within the crowd, agents effectively communicate with those behind them, reducing the pressure below χ . This prevents

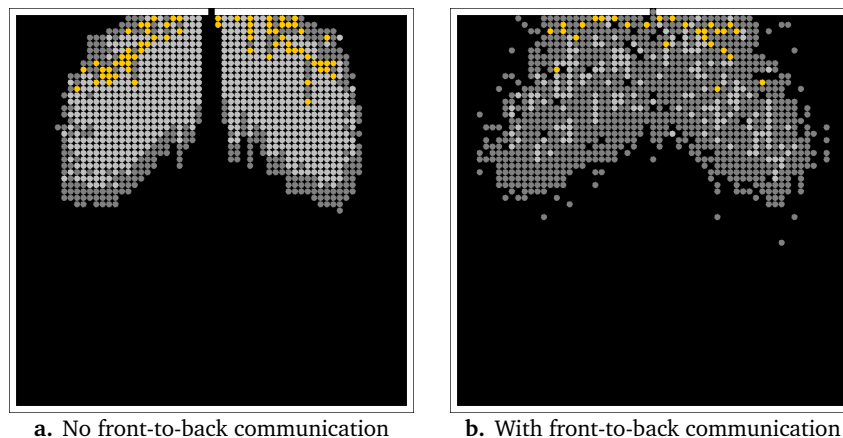


Figure 28: Disappearance of the aisle region with front-to-back communication. Light grey agents exceed χ and are pinned in place. Yellow dots indicate injured agents. Both with $k_S = 10$, $k_D = 0$, $\phi = 105$, time step = 175.

aisle creation as agents control the force upon them sufficiently to allow the lateral steps that they would otherwise be prevented from making. With no force breaks required, high ϕ values (where injuries drop to zero) do not impede the exit rate and the aisle does not form (see figure 28b). Low ϕ values continue to produce high numbers of injuries and a consequent decrease in exit rate through clogging of the exit area by injuries.

This suggests that there may be multiple mechanisms at work to prevent the formation of the aisle region in crowds. We have now suggested three mechanisms that reduce aisle effects, namely introduction of injuries, allowance of diagonal steps (proposed as future work), and now front-to-back communication. The fact that front-to-back communication reduces injuries and eliminates the formation of the aisle indicates its effectiveness at controlling force in the model, as well as emergent crowd structures that depend on force.

In chapter 6 we used the protective effect of injuries to argue that force in crowds was an essential aspect of crowd modelling, particularly at high densities. The results of this section, although they simplify the high density behaviour, do not invalidate our point. The low ϕ results indicate that when injuries occur, the non-linear response in the exit rate persists.

The scenario of section 8.4, just discussed, shows that under laboratory-type conditions front-to-back communication can indeed reduce injuries at the cost of a longer exit time. At high density and high drive, the most dangerous crowd scenario, injuries are reduced and exits are facilitated. We may well ask how front-to-back communication affects a more realistic scenario. We turn to that question now.

8.5 Who concert disaster scenario

In order to provide further support for the hypothesis that front-to-back communication can improve crowd dynamics, we wanted to evaluate the effect that front-to-back communication could have in a non-laboratory situation. We chose to simulate the Riverfront Coliseum concert disaster because of an encouraging and detailed account of communication and social behaviour (Johnson, 1987b), and the good descriptions of the physical surroundings and events (Johnson, 1987b; Chertkoff & Kushigian, 2001).

8.5.1 Reconstructing the plaza at Riverfront Coliseum

The plaza of the coliseum has been altered since the time of the disaster; we attempted, however, to reconstruct its physical dimensions based on three sources: a *news photo* taken by the Cincinnati Enquirer the night of the disaster (Cincinnati Enquirer, 1979/1999), an *aerial photo* of the original plaza and building ([Riverfront Coliseum], n.d.) and a *schematic* from Johnson's account of the disaster (1987b). Johnson's schematic is reprinted here as figure 29 (to which all reference marks refer).² Despite perspective effects, we were able to produce a consistent plaza geometry.

The news photo shows a bank of 8 doors at location A. We supposed these doors were full size (914 mm) establishing our scale. The news photo also shows a large squarish column (~ 4 doors wide) at i, and a smaller column (~ 1 door wide) at the building-lobby junction near A (call this BLJ). These estimates fix the length of wall A at 11.8 m. The view of the doors at B in the news photo is partially obstructed, but Cherktoff and Kushigian say there were 8 doors there as well. From the news photo it appears the doors at B were flanked by column i and an identical column on the opposite side, setting wall B at 14.6 m long. The wall measures set the scale for figure 29 (used for all front-to-back estimations, scale set from length of A) and the aerial photo (used for all side-to-side estimations, scale set from the doors at B). The model's x axis is parallel to wall B; the model's y axis is parallel to wall A.

The lookout at the right of the plaza is an isosceles trapezoid with height 3.9 m, whose parallel sides measure 37.1 m and 44.6 m. Remaining side-to-side dimensions (all projected onto the x axis) are: BLJ to E 17.7 m, E to ii 4.9 m, ramp width 15.6 m. Remaining front-to-back dimensions (all projected onto the y axis) are: ii to F 17.8 m, BLJ to iii 47.5 m, E to BLJ 10.9 m, the wall of C opposite A is 5.9 m.

We needed to set the size of a grid cell in the floor field model, choosing 56 cm, the commonly cited (if dated) reference human body width (Pauls, Fruin, & Zupan, 2007). The above estimates were quite consistent, only requiring minor adjustments of ± 1 cell to make the

²The other two photos cannot be included for copyright reasons. The Cincinnati Enquirer refused to give permission to reprint their photo. The aerial photo appears unattributed on the Internet; inquiries made as to the source of the image were not returned. The reader is directed to the URL addresses in the references if he or she wishes to review the images.




Figure 29: Plaza at Riverfront Coliseum, reprinted from Johnson, 1987b, p. 366, fig. 1. The shaded zone indicates the area of highest pedestrian density. Johnson's reference marks are in upper case: A-South doors, B-North doors, C-common lobby, D-casualties, E-broken door, F-officers. Our reference marks have been added in lower case, and are further identified in the text.

various structures line up correctly. The 56 cm cell implies approximately 13 cells per bank of 8 doors. We interpreted remarks that insufficient doors were open by supposing that half of the doors were open, so we have included 6 exit cells in the centre of each bank.

Although a police officer inside the coliseum reported “there must be 8000 people standing on the outside trying to get in,” (qtd. in Johnson, 1987b) we feel from the description of the officer’s circumstances that this was conjecture rather than a reliable estimation. Given our estimates and our square grid cells, 5000 agents provided a very high density sufficient for our purposes. We initially distributed these agents randomly within the shaded (high density) area of figure 29.

8.5.2 Analysis of the disaster

We ran the simulation using the same settings used in the laboratory experiments, except that we fixed k_S and k_D at 1.0 and 0.5 respectively due to the high forces produced by this scenario (see below). We sampled a range of ϕ values from 100 to 300, with and without the front-to-back communication simulation active, for the first 350 time steps of model execution. As with previous sections, we report on means of 50 simulations for each parameter combination.

First we consider the scenario without front-to-back communication enabled. We found that forces were high and the scenario was prone to producing injuries. In such a large scenario, injury thresholds have a strong effect on injury distribution. This is because, at low values of ϕ , agents rapidly become injured throughout the crowd due to the relatively small numbers of other agents required to inflict damage (see figure 30a). Injured agents throughout the crowd then prevent force from propagating over long distances. When ϕ is higher, injuries are more focused because many ranks of agents are required to generate the cumulative forces required to cause injuries.

In observing the simulation, we noted that injuries tend to appear first at a force hot spot. While conscious of the fact that our model is only a qualitative simulation and reconstruction from secondary sources, our observations of the simulation suggest that the area around D (where injuries occurred in the real disaster) is particularly prone to being a site for injuries




Figure 30: Riverfront Coliseum scenario without front-to-back communication. Exit cells are green, agents are grey, agents experiencing force above χ are light grey, injured agents are yellow.

(figure 30b). When these forces continue to be applied at a high level, our model shows injuries continuing to propagate from the area around D onto the plaza, toward the ramp (figure 30c). We shall discuss the formation of the hot spot and the reasons for the continuing injury propagation separately.

Hot spot formation. Our model is most sensitive to injuries at location D, where injuries did occur on December 3, 1979. Fruin proposed an explanation for this location of injuries at Riverfront Coliseum, suggesting that “crowd pressures probably came from both directions as rear ranks pressed forward and front ranks pushed off the wall” (1993). Our observations of the model point to a different mechanism for the force hot spot at location D. In our simulation, the hot spot is located at the crossing point of two sources of force (figure 31). When the geometry and lack of intermediate force breaks allowed, pushing and leaning forces originated from those on and exiting the ramp, quickly overwhelming the free choice of agents in front of them on the plaza. Agents on the plaza were forced to add their pushing and leaning forces toward D

rather than allowing them to aim for B. Thus, large cumulative forces travelled over significant distances from the ramp toward D. These forces were compounded by a second source of force, as agents near the coliseum end of the lookout generated additional perpendicular forces toward D.

Compared with Fruin's, our hot spot theory gives a more plausible explanation for these injuries. The hot spot theory accounts for the fact that injuries occurred at a specific point: the crossing point of the forces. This contrasts with Fruin's suggestion which is more compatible with a line of injuries between A and F. Our reconstruction of the coliseum also suggests a reason that injuries were set back from the door at D, rather than occurring right at A; the geometry of the ramp is such that the door area is not accessible on a straight line from the ramp (see figure 31).

Continuing injury propagation. The continuing propagation of injuries onto the plaza (figure 30c) was not reported in the disaster. We note here three possible reasons for this discrepancy. First, our model does not distinguish between fatal and non-fatal injuries; non-fatal injuries in the crowd at this event are poorly documented and there may indeed have been high forces propagating to these locations. Second, other non-modelled factors in the scenario (e.g. the sudden increase in drive toward the exits brought on by the band's warm up) may be relevant in limiting peak injuring forces to particular time periods (a point we shall return to in section 9.1.3). Third, it may simply be that our abstract force model based on a uniform grid of cells lacks the fidelity to model additional dissipative factors that prevented injury in these locations despite continuous pushing.

8.5.3 Comparison with front-to-back communication

We turn now to the results (figure 32) of enabling the front-to-back communication simulation in the Riverfront Coliseum scenario. We found that, as with the laboratory scenario, agents were able to control forces over a long distance, substantially reducing the number of injuries. As the lobby was the movement goal for agents, it is not surprising that at high values of ϕ (when injuries are low) remaining injuries occur in this area. Observations of the model suggest




Figure 31: Perpendicular forces (*white arrows*) create injury hot spot (*grey circle*) at Riverfront Coliseum. Note that hot spot cannot move to the left toward the doors, as ramp wall constrains further leftward movement.




Figure 32: Effect of front-to-back communication at Riverfront Coliseum

injuries were more distributed, with hot spots for injuries near i , near the exits and between A and E. We found that injuries within the model did not have an impact on egress rate. As ϕ increased, making agents less susceptible to injuries, exit rates remained constant.

In summary, introducing front-to-back communication into the model resulted in a more patient, less forceful crowd. This had a consequent effect to decrease the exit rate, which was also unaffected by numbers of injuries.

8.6 Return to Lenin Stadium

In section 7.4.2 we generated a scenario similar to the one that occurred at Lenin Stadium (section 2.2.2). At Lenin Stadium, freezing-by-heating caused deaths and injuries when two groups of people attempted to pass through one another at high drive. Although we did not use the physical dimensions of the real event, we did reproduce a barrier formation similar to the one that led to the deaths in that situation. In our reproduction, the barrier arose because of disagreement among agents about spatial facts. When agents were able to communicate, the barrier formation was dissolved because agents resolved their spatial disagreements.

At the end of chapter 7, we raised the question of whether the spatial communication of chapter 7 could have helped at Lenin Stadium. We concluded that as spatial disagreement was not present, the spatial communication of chapter 7 was the wrong tool for the job in that case.

In the present chapter, however, we have introduced a different type of communication aimed at reducing force and injuries that are due simply to pushing forces originating at the back of a crowd, and this is more suitable for the scenario of Lenin stadium. Would improved front-to-back communication (as opposed to the spatial communication of chapter 7) have helped at Lenin Stadium? Recognising that — in the absence of an accurate physical simulation — our scenario is highly abstract, we can nonetheless enable front-to-back communication in the scenario of section 7.4.2. For comparative purposes we have used only the high drive parameter values from section 7.4.2 ($k_S = 7$ and $k_D = 0$).

Although we have just argued that the front-to-back communication condition is the one most applicable to Lenin Stadium, it is not necessarily the only option in other scenarios where

barrier formations may arise. Barrier formations generally are also worthy of study. This return to the scenario of section 7.4.2 also offers the opportunity to study a combination of both front-to-back and spatial communication on barrier formations without regard to Lenin Stadium. Accordingly we present this additional data as well.

8.6.1 Results

Injury results, shown in figure 33, suggest that front-to-back communication alone (our Lenin Stadium case) is successful at moderating force in a barrier formation. Accordingly, we conclude that improved front-to-back communication at a disaster like the Lenin Stadium disaster could have been protective.

As discussed, although spatial communication might not have helped at Lenin Stadium, this does not mean that no situation featuring a barrier formation could benefit from spatial

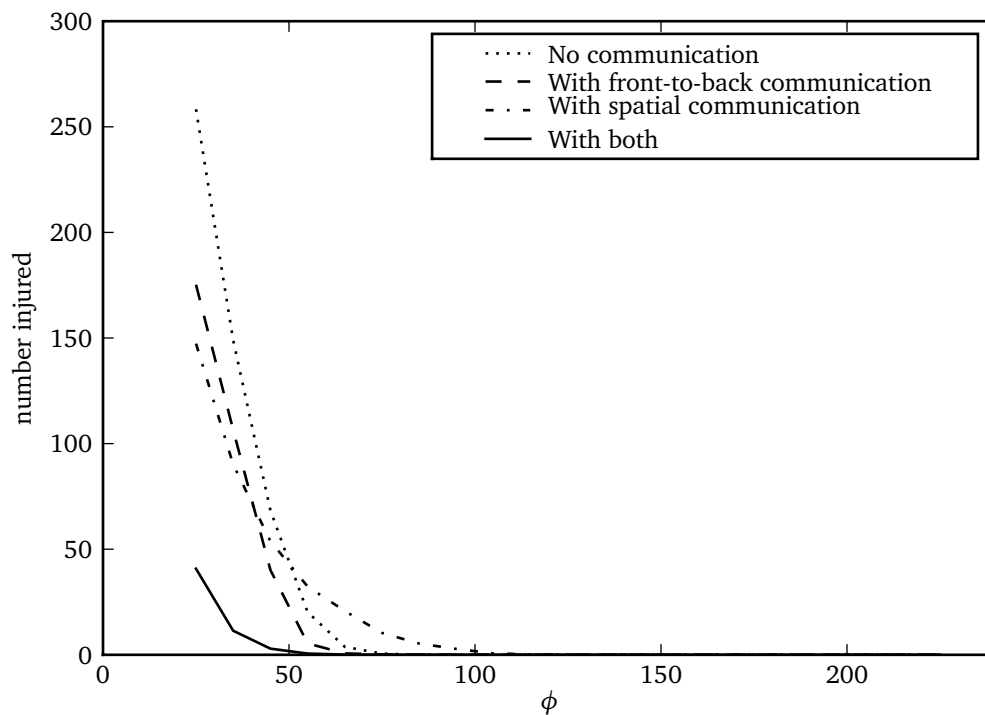


Figure 33: Front-to-back and spatial communication in a barrier formation

communication. We have characterised spatial communication alone in section 7.4.3, but what if front-to-back communication was also present in such a scenario?

When we enabled both types of communication we obtained the best injury performance, with spatial communication breaking up the barrier formation and front-to-back communication protecting agents near the exits. The x -intercepts of the curves also are revealing because they identify the peak force experienced by agents within the crowd in the various cases. As discussed in section 7.4.3, spatial communication generates higher forces because all agents quickly learn the location of the true exits, rush to the doors with high drive, and forces can be high due to higher densities under these conditions. Adding front-to-back communication to the spatial communication was protective in this regard, as peak forces were lowest in this condition as read from the x -intercept of figure 33. This is further evidence of the ability of front-to-back communication to mitigate the forces generated at high densities and high drive.

Turning to exit results, we first note that as with most of the results of chapter 7, the exit rate was not affected by injuries. Accordingly, results are shown in figure 34 to allow for comparison between ψ and comparison of communication types. Introducing front-to-back communication alone (despite the results of figure 33 showing reduced injuries) did not allow for many more knowledgeable agents to exit in the time allotted; although it reduces injuries by limiting

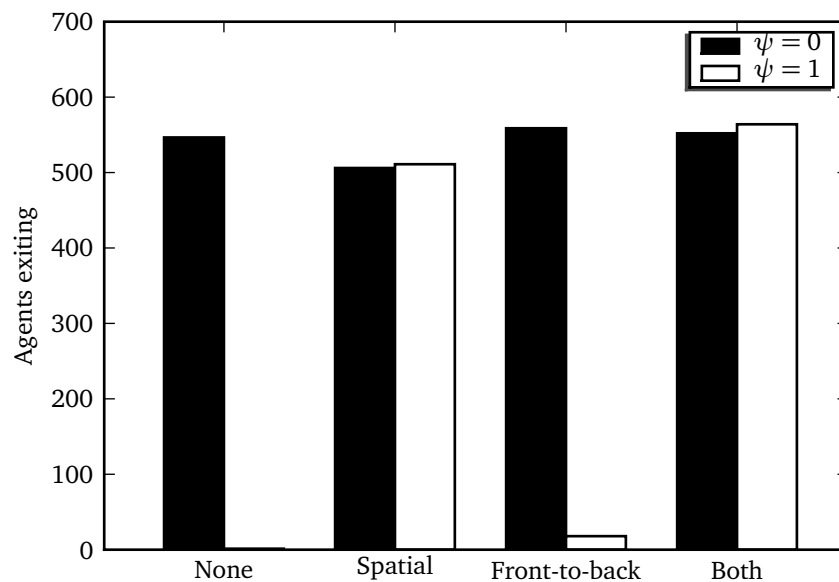


Figure 34: Comparison of communication styles on a barrier formation. $\phi = 125$, $k_S = 7$, $k_D = 0$

pushing, front-to-back communication also eliminates drive to exit by setting $m_S = 0$. (In other words, it does not dissolve the barrier formation, it simply reduces the danger inherent in being a part of it.) In combination with spatial communication, however, front-to-back communication dramatically improved exit performance. This is only partially evident in figure 34 due to a ceiling effect in which all agents were able to exit in each replication of this parameter combination. The mean time to exit was 146.2 time steps. This rapid exit was possible because, as discussed above, front-to-back communication prevents forces from building up and limiting movement at high drive.

8.7 Model validity

As we have done for the previous two demonstrations, we now turn to the question of the validity of our front-to-back communication model. This demonstration differs from the previous two, as our goal is to demonstrate the methodology at work in a more speculative mode. The properties of front-to-back communication behaviour have not previously been described formally, and indeed front-to-back communication itself has been defined and discussed only in the context of its absence from certain crowd disasters. Accordingly, providing an argument for our model's relevance to the real world is very important.

Implementation. The validation goal at the implementation level is to determine if the microscopic human factors have been appropriately implemented. This includes the simple question of implementation correctness, but also the appropriateness of the formalism selected for implementation. The reduction for front-to-back communication is quite naturally implemented in the context of the floor field model, which has proven to be an excellent formalism for all of the demonstrations of the methodology. Through observation of the model we have verified that the four basic processes of initiation, retransmission, action and decay correctly implement the reduction. No simplifying assumptions were required.

Reduction. At the reduction level, the task is to determine whether the reduction is an appropriate abstraction of the specification. As discussed above, we have taken quite an abstract view

of our specification. According to the specification, cues for initiating front-to-back communication and for determining the back of a crowd are complex; the reduction provides for simple force-based mechanisms in both cases. These mechanisms, however, are seen as conservative in the sense that the omitted complexity would likely result in easier triggering, and more accurate retransmission. It may be possible to introduce more complexity into the microscopic model to capture some of the nuances introduced by these other cues, however this is left to future work.

One advantage of implementing and testing a model for a behaviour that is not completely understood is that, in so doing, we ensure that the behaviour is completely described. For example, problems determining the location of the back of a crowd might not be anticipated if it were not required to provide a complete abstract model of the behaviour. The results of the simulations demonstrate that the abstract model is able to reduce injuries and manage pushing forces. In this respect, the reduction has demonstrated its sufficiency as a mechanism for distributing a signal from the front to the rear of a crowd, with increased safety in mind. It is able to meet the goals of the specification, and as all the major elements of the specification are represented in the reduction, we conclude the reduction is acceptable to meet our specification of front-to-back communication in crowds.

Specification. The purpose of validation at this level is to determine the extent to which the model is a representation of the behaviour of real crowds. As discussed in section 4.3.2, abstract microscopic models can provide a forum to explore how things *can* work, and to explore mechanisms. Like in many areas of research that seek to explain complex human behaviours, the results of our analyses should be viewed as one piece of evidence out of many; we begin to develop a coherent picture of a complex behaviour as other disciplines join in and we accumulate multiple pieces of confirming evidence that reinforce our theory about how the behaviour works.

This chapter demonstrates the methodology in a prospective mode to investigate an emergent crowd behaviour that is not well understood. Given the prospective nature of this chapter, there are few results from other disciplines to mesh our results with. Although we can appeal to case studies and Fruin's and Pauls' general comments, it is fair to say that our model does not prove that front-to-back communication exists, or that it works in real crowds. We can say, however,

that if front-to-back communication works as our specification suggests, it can in principle reduce injuries in circumstances of extreme force and high density and drive. This is more than could be said before.

Our hypothetical specification of front-to-back communication was inspired by the three case studies in chapter 2. We have considered two case studies in this chapter so far. One in depth (Riverfront Coliseum) and one in a more abstract manner (Lenin Stadium). In the case of Riverfront Coliseum, we argued for a similar pattern of injuries between the model results and published accounts of the disaster; that these injuries were dramatically reduced by introduction of front-to-back communication supports our view of this kind of communication as protective. At Lenin stadium, a barrier formation developed as two groups tried to move through each other at high pressure; in our model front-to-back communication reduced the injuries (although it did not eliminate the barrier). Finally, we are left with Fruin's comments that crowd disasters can be understood as a failure of front-to-back communication. Our results support Fruin's point in that switching off front-to-back communication can have serious consequences in the case of high density, high pressure crowds (see section 9.1.3 for more on this point).

8.8 Discussion

The work of chapters 6–8 suggests a crucial role for people at the back of a crowd at a bottleneck (such as in egress or ingress). These people, although far from doors, can nonetheless constrain the behaviour of those near the exit, jam intervening people in place, and ultimately induce injuries by activating and targeting the leaning forces of others. These injuries may occur far from the rear agents; as we have seen from Johnson's account, however, it is likely the media will lay the blame on those closer, who actually had no control over the outcome.

The introduction of front-to-back communication places an extra responsibility on those at the rear. They are the only ones who can actually choose to withhold pushing forces, and hence prevent the cascade of pushing and leaning that causes injuries, and compromises exit rates through jamming. With front-to-back communication, those at the rear, in principle, can dynamically reduce drive as needed to maintain the safety of others, and this mitigates some

danger in a high-density, high-drive crowd situation. If front-to-back communication is functional and honoured then this produces the best exit rates. However, if front-to-back communication is not honoured, then, as seen in figure 27, high drive is counter-productive, producing many injuries at high density.

It is one thing to say that honouring front-to-back communication produces the best exit rates, and another to say there will be time for everyone to exit. When time constraints become extreme (e.g. in fire egress, in crazes or even in the case of dwindling time to enter a concert), the rearmost crowd members are the ones who have the least probability of reaching the exit in orderly egress. Competition can develop, and a situation recalling the prisoner's dilemma (e.g. Axelrod, 1984) can emerge. A globally optimal solution is for the rearmost agents to co-operate fully in front-to-back communication, increasing patience and ensuring the best exit rate for those in front. From the point of view of the individual in the rear ranks, however, this does not maximise personal likelihood to exit; rather than increasing patience, changing positions in the crowd (to move closer to the exit relative to other agents) would increase the probability of a successful outcome. In this situation, the rearmost individuals may well become less willing to honour the front-to-back communication signal, particularly if they do not understand the consequences of this behaviour (see also section 9.1.1). Regardless of the outcome of this ethical dilemma, if forces *are* applied from the rear the result is the jamming and injury reported in chapter 6; rearmost agents, along with those in front, are negatively impacted by the faster-is-slower paradox in this case.

All of this suggests investigating the effects of less-than-perfect compliance with the front-to-back communication protocol. Although we leave it to future work, adding probabilities of initiation, of honouring the front-to-back communication protocol and of re-transmitting the signal could easily be added to algorithm 4. (The probability for initiation can be added to line 5 of procedure `PREPARE-AND-INITIATE SIGNALS`; for honouring to the start of the `HEAR-SIGNAL` procedure; for re-transmitting to line 17 of procedure `STEP`, with a few minor housekeeping details in the state.)

8.9 Summary

Although previous research has identified failure of front-to-back communication as an important characteristic of crowd disasters, it seems fair to point out that without examples and a description of successful front-to-back communication it is hard to put its failures into context. In other words, in order to speak of a behaviour that has failed, we must be able to say *how* the failure is different from successful cases of the behaviour.

Our goal was to study this phenomenon within a microscopic model to determine what potential benefits front-to-back communication may provide, and the effect of introducing front-to-back communication into a crowd situation. We have proposed a model of front-to-back communication that includes: initiation based on a loss of individual movement control, retransmission through successive dyads, behaviour modification for parties to the communication in the form of reduced voluntary pushing and movement drive, and a time-based method of decay of this behaviour modification. We implemented this conception of front-to-back communication in our modified floor field model, building on the product of chapters 6 and 7. Our results suggest that front-to-back communication does have the potential to reduce damaging forces in a laboratory scenario, at the cost of reducing the exit rate due to increased patience. At high drive, front-to-back communication decreased injuries and increased exits.

To evaluate front-to-back communication in the context of a more realistic scenario, we reconstructed the geometry of the plaza at Riverfront Coliseum in Cincinnati, Ohio that witnessed the Who concert disaster of December 1979. We based our approximation of this space on secondary sources, such as published diagrams and photos, and obtained a consistent conformation. Although our simulation is designed to obtain qualitative results relating to the introduction of front-to-back communication, we found that our model predicted a force hot spot, with subsequent agent injuries, in the same area that saw real injuries in 1979. The results suggest that forces built up over a long distance from the plaza ramp toward the coliseum, and were compounded by perpendicular forces from the plaza lookout; this improves on Fruin's suggestion that injuries occurred when forces from people pressing off the front wall met forces coming from the rear. When we added front-to-back communication the number of injuries decreased

substantially, as did the exit rate. These results suggest that front-to-back communication may be protective in crowd situations, at the cost of exit rate.

We have suggested three areas for future work on the front-to-back communication simulation. First, taking non-force cues into account in the initiation and retransmission of the front-to-back communication signal. Second, investigating initiation, co-operation and retransmission probabilities to establish the degree of compliance required for benefits to occur, and whether benefits degrade smoothly with a decrease in compliance. Third, investigating the change in effectiveness brought on by changes in decay probabilities. We further note that unlike at the Who concert tragedy, many crowd venues now have voice communication systems and overview facilities for trained crowd management personnel. We would like to consider the interactions that global information and communication may have with front-to-back communication in a crowd situation.

We have now demonstrated the Microscopic Human Factors methodology in three different applications. From these three applications, however, come certain joint conclusions, to which we now turn. We shall then comment on the value of the methodology in general and return to a consideration of Sime's concerns.

Chapter 9

General discussion

The work of chapters 6–8 was concerned with demonstrating the use of the Microscopic Human Factors methodology in adding human behaviour to microscopic models. Each chapter has built on those that came before, and each has come to certain conclusions about crowd behaviour along the way. When taken in concert, however, they have more to tell us, particularly when combined with important papers from the literature. The first purpose of this short chapter, then, is to present this general discussion of model results, and we shall do so in section 9.1. These demonstration chapters, however, were presented first and foremost as examples of the Microscopic Human Factors methodology at work. We shall turn to a discussion of the methodology in section 9.2. The methodology itself was presented as a way to bring Sime’s concerns into the microscopic modelling domain, and we shall return to these questions in section 9.3. Thus, this chapter unwinds the themes of this thesis in reverse order, bringing us back to our goals and the motivations of chapter 1.

9.1 Crowd disasters

9.1.1 Training in front-to-back communication

In section 8.8 we discussed the important role of those in the rear who, our work suggests, are the only ones with the freedom of action needed to prevent jamming and injuries in crowds. We pointed out that rearmost agents — by pushing — unleash and target the leaning forces of others in a high-density crowd. But, do people in crowds understand this role? If we take media

reports to represent general understanding then they do not — as evidenced by the focus of the media on those near injuries who are literally forced to injure others (Johnson, 1987b).

People in emergencies take the actions that make sense to them at the time, given their prior knowledge and experience, the information available at the time and available cognitive processing resources under pressure (Proulx, 2002). For example, to an untrained person, breaking a window to vent room smoke in a high-rise fire may make perfect sense; someone with knowledge of the stack effect and building airflow might make a different decision given their increased knowledge and the circumstances. Similarly, in a crowd disaster, it does not seem likely that most crowd members will recognise the counter-intuitive faster-is-slower paradox, or understand the important role of those at the rear. Indeed, if a disaster like Riverfront Coliseum does occur, because of the low awareness of the mechanisms involved, those at the rear may never realise the part they played in it. In these circumstances, we cannot fault people at the back who engage in what they believe to be harmless pushing without recognising the danger, even if they are the cause of forces and injuries (see also sections 2.3 and 2.4).

If the people involved in crowd events are not at fault but crowd disasters continue to occur then we must determine how these situations can best be prevented in the future. One way may be to directly educate people about the dangers of high-density crowds, the important role of those at the rear and the value of front-to-back communication. This would fall under the category of altering crowd members' prior knowledge and experience. For example, a social convention for an easily recognised signal that can be propagated throughout a crowd could be devised.

Before recommending formal instruction in front-to-back communication for all dangerous crowds, we must consider whether all crowd disasters are well suited to improved front-to-back communication. Guided by Fruin's statement that "most crowd incidents exhibit a failure of front-to-back communication" (1993), our study suggests that if front-to-back communication is present in crowd situations it improves safety. However, this does not mean that it *can* be introduced into *every* crowd situation. In particular, some disasters happen very quickly: at Bethnal Green, for example, "in a matter of seconds there was built an immovable and interlaced mass of bodies 5 or 6 or more deep against which the people above and on the stairs continued

to be forced by the pressure from behind” (Dunne, 1943/1999). It is hard to see how even perfect training in front-to-back communication could have produced measurable results in this time-frame, given the number of ranks built up and waiting to enter.

Although it is hard to see how the events of Bethnal Green could have been slowed down, the same is not true of all crowd disasters. While some crowd situations unfold very rapidly, there are crowd situations that entail considerable risk and at which crowd forces build up over a much longer period of time. Large, slow, organised crowd events (e.g. the Hajj) may support instruction and training in a front-to-back communication protocol. We are aware of one such protocol that we call the Halifax fire drill protocol, taught to students as young as 8 years old: when following a line of people to an exit during a smoky fire, the turning around of the person in front is a signal that the exit is blocked.¹ Our results suggest this strategy of formalising simple front-to-back communication protocols is promising and merits further study.

The concerns about timing discussed in this section lead to a suggestion for future work: In our simulation the speed of force, information and movement is the same (all have a rate of 1 floor field cell per time step); an abstract investigation of the sensitivity of egress rate and injuries as measures of communication effectiveness when varying force and information propagation speed would be quite interesting.

9.1.2 Often there is plenty of time to support the outcomes we want

Directly training people in front-to-back communication is not the only way to go about improving crowd safety. We have raised the question of time in the previous section; our results directly support a view of crowd safety that has more to do with time than communication.

As we have seen, agent drive has a large impact on crowd safety. The provision of timely information can be a factor in determining agent drive. Speaking specifically about fire egress, Sime has noted (e.g. 1995) that the egress time is composed of pre-movement components t_1

¹This protocol was related to us by a former student of the Halifax school board (J. A. Schellinck, pers. comm., 2007). Enquiries at the school board (D. Bell, pers. comm., Aug. 22, 2007) and at the Halifax fire department (C. MacDonald, pers. comm., Aug. 17, 2007) indicate that this policy was not implemented board-wide and a fire department officer with 25 years of experience (11 years in fire prevention) was unaware of it. Although it cannot be conclusively shown this is the case, this protocol was likely the local creation of a staff member in the specific school, and has not been formally studied.

(such as detecting fire cues, deciding to leave and preparing to leave), and movement time t_2 . Total egress time t is thus given by $t = t_1 + t_2$ and not by $t = t_2$ as generally modelled (a limitation of our model as well).

In part, these comments are due to work by Proulx and Sime who noted that early, accurate, specific information allows people time to leave a building in an orderly manner (Proulx & Sime, 1991). In other words, it is possible to ‘buy time’ by starting egress sooner; often time allows orderly egress, but egress is delayed due to a late warning (sometimes for fear of causing ‘panic’) or due to insufficient drive to exit. Proulx and Sime argue that “rather than avoiding ‘panic,’ a delay in warning the public is self-fulfilling. When an incident such as a fire gets out of hand, the delay is paradoxically a major determinant of flight behaviour, crushing and deaths in major crowd disasters” (1991).

Let us consider Proulx and Sime’s statement in the context of our simulation. We have studied three cases across chapters 6–8: high drive, medium drive and low drive. At low drive, agents do not move to the exits in large numbers. When a limited number of time steps is available to them, many of them fail to exit due to this low drive. This accords with Proulx and Sime’s views: in emergency egress low drive wastes time. As time becomes a factor, Proulx and Sime describe a transition to flight behaviour, which we interpret to correspond to high drive. The results of chapter 6 suggest this is a dangerous regime. Despite this, chapter 8 (if front-to-back communication works as our specification suggests) offers the hope that, if rigorous adherence to front-to-back communication protocols is observed, injuries can be prevented and high exit rates can be maintained. As we have argued in section 8.8, however, the people at the rear may be boxed into an impossible position: push and cause a jam through faster-is-slower, or co-operate and risk being trapped as time runs out. Clearly the better solution is the medium drive case; if time allows (and this requires prompt notification to building occupants who must then respond promptly and seriously) then the medium drive case allows an optimal number of people to exit safely.

Could the medium drive case have been encouraged at Riverfront Coliseum? The work of Proulx and Sime suggests an important role for better information. If a public address system had been available, for example, the promoters could have informed fans that the band was warming

up, that the doors would open behind schedule, that the concert would be delayed to allow everyone to enter, and that there was no need to rush. Of course, our model does not tell us that this specific suggestion would have worked, but rather that some method that produced medium drive would have been a better bet than hoping for front-to-back communication. This could have included pre-event factors, such as proper event planning, facility capacities, emergency plans, crowd management, etc.

This discussion suggests two pieces of future work: investigating dynamically changing agent drive, and modelling pre-movement time. If some fire disasters are characterised by an initially apathetic response, followed by a high-drive phase, it should be possible to investigate this within the model by providing for agent-by-agent dynamic k_S and k_D values. Changes in these parameters could be driven by time, or by discovery of fire cues along the lines of chapter 7. Modelling of pre-movement time is a task as daunting as it is needed, as it requires modelling decision-making processes that are totally dependent on information in the shape of fire cues, inter-agent milling and communication. Despite the challenge, decision-making under pressure (including the role of emotion) is a fertile area for research in cognitive science, and one that could produce significant benefit to crowd and fire egress modelling.

9.1.3 Was it a failure of front-to-back communication?

Let us now return to Fruin's statement that "most crowd incidents exhibit a failure of front-to-back communication" (1993). Is this statement supported by our research? In a strictly literal sense the answer is yes. After all, Johnson's research (1987b) suggests that people are at their best in disasters, and if those at the back of a crowd understood the direct consequences of their actions we have no doubt that crowd disasters would be much rarer. Also, we have shown that when front-to-back communication is functioning crowd safety is enhanced, while in the absence of front-to-back communication injuries increase. The work discussed in section 9.1.2, however, gives us an alternate explanation that is equally supported by our model: The reason crowds find themselves in need of enhanced front-to-back communication is a transition in agent drive from low to high as time dwindles and competition appears.

In our view, the more nuanced answer to the question of this section is that Fruin's statement is a consequence of a broader view of Proulx and Sime's work. This view depends on generalising Proulx and Sime's observations to low → high drive transitions more generally, dropping their emphasis on information management and fire egress. Let us consider two of our case studies in this light: At Bethnal Green (section 2.2.1) what was an orderly (moderate-drive, we suppose) descent into an air raid shelter became a high-drive flight when the noises of war became apparent, prompting the pressure on the stairs, the subsequent fall, and the jamming transition that ensued. At Riverfront Coliseum, Chertkoff and Kushigian describe the surge in the crowd (transition from stasis → high-drive) that occurred as the doors opened, leading to a fall due to uncontrollable forces at the front.

Elements of the descriptions of these events (Dunne, 1943/1999; Johnson, 1987b) make clear that failure of front-to-back communication did occur in both of these cases at the time of transition to high drive. If front-to-back communication along the lines of our hypothesised mechanism had been successful, our analysis suggest it might well have reduced injuries. In this sense, our work supports the idea that a failure of front-to-back communication was exhibited. However, without the transition in drive it is not clear that either event would have become a crowd disaster. In our view, it may be more promising to look toward drive in characterising crowd disasters.

Although precise timing for these disasters is not available, it is interesting that both featured salient auditory environmental cues that could have acted to synchronise the transition from low to high drive throughout the crowd. Namely anti-aircraft guns at Bethnal Green and the sound of the band warming up at Riverfront Coliseum. This may be relevant given our hypotheses about time to propagate information and force from section 9.1.1. Our conception of front-to-back communication depends on a slow build-up of force that allows time for initiation, for information to travel backwards and for pushing to discontinue; to avoid injury this process must run to completion before force builds to the threshold for injury near the front. If a synchronised transition to high drive occurs in a high-density crowd, however, high forces may occur before the front-to-back communication process runs to completion. On this view, injuries could have been mediated by effective front-to-back communication, except that conditions for effective

front-to-back communication are absent. Accordingly, if a synchronised transition from low to high drive did occur, it may be a better explanation to say that the co-ordinated transition in drive caused the disaster.²

The discussion of this section raises three questions that could be fruitfully answered in future work. First, a review of crowd disasters could be undertaken, looking for the type of synchronisation that we hypothesise could have been at work at Bethnal Green and Riverfront Coliseum. Second, descriptions of both of these disasters (and of many others besides) discuss the fall of an agent as the seed for jamming and injury; it would be interesting to return to the work of chapter 6, raising ϕ to a threshold that results in no injuries, but then artificially inducing a seed injury in an effort to understand its effects on the rest of the crowd, and why falls are so salient in crowd disasters. Lastly, we have suggested (in section 9.1.2) an investigation of dynamic drive changes in crowds, but investigating the effects of synchronisation of those changes would be a valuable addition to this future study.

9.2 The Microscopic Human Factors methodology

The studies of force, spatial information discovery and communication and front-to-back communication have each been generative in their own right, and as we have just seen, have provided for some interesting discussion in concert. Their larger function in this thesis, however, is as demonstrations of the MHF methodology. In this section we evaluate the methodology itself, the use of which we found to have several benefits in the demonstration chapters.

9.2.1 The value of the specification

An important benefit of the methodology was that it required the development of a specification based on what we know about the crowd phenomenon at issue; it required us to make explicit our hypotheses about those aspects of the phenomenon that are not well understood. The

²An intuition pump may solidify this view. Given the auditory stimulus present in both cases, if we suppose that both the Riverfront Coliseum and Bethnal Green crowds were wearing earplugs and did not hear the concert beginning or the noise of war, it is not difficult to imagine that the drive transition might not have occurred. Given continuing orderly progress there would be no impetus for a fall, and likely no disaster.

specification must be centred on crowd behaviour, and this ensured that each demonstration began with a discussion of crowd behaviour and its mechanisms. Beginning with this discussion is very important as it places the focus on the mechanisms of real crowd behaviour, and encourages an abstraction that preserves these behavioural mechanisms. A less-desirable alternative might begin with an argument for a modelling formalism like particle dynamics, justifying it by an inappropriate analogy (e.g. that people are like seeds in a hopper) and proceeding to model ball-bearings instead of people (see sections 1.1.2 and 3.2 for more on this point). The MHF requirements help us to guard against such a strategy; when we begin by describing people, for example, it becomes quickly apparent that seeds in a hopper are a poor metaphor for people in crowds.

9.2.2 The value of the reduction

In our experience using the methodology, the reduction level is the locus of the most important modelling contribution. The reduction layer makes explicit — separate from any questions of modelling practicality or implementation formalism — the known or hypothesised essential aspects of a behaviour or crowd phenomenon. This, in our view, is the primary theoretical output of a model that exists to explore the mechanisms at work in crowds; it is the source of the explanatory power in a microscopic simulation. The absence of an explicit reduction in a microscopic model leaves to the reader the critical task of extracting the authors' hypotheses about how the world works at the essential and abstract level consistent with the microscopic model's implementation. Leaving to the reader the responsibility to imagine how a real crowd might be similar to the implementation described is leaving an important step in the validity of the model undefined; if it is not clear how the implementation relates to the world, then it is not clear how the results of the implementation do either.

We were surprised that other models (e.g. Burstedde et al., 2001; Helbing et al., 2000b) do not appear to provide this abstract description of the mechanism at work in crowds, separate from their implementation and formalism. These models do not make clear what abstractions in the model are made for the sake of the formalism, and what abstractions are considered to be

theoretical positions concerning mechanisms in crowds. By contrast, our demonstrations show this important task is accomplished if the methodology is followed. In each reduction (sections 6.2, 7.2 and 8.2) we outlined the critical processes that, in our view, must be present in order for the behaviours to happen; these represented our hypotheses of how the world essentially works. In each of sections 6.9, 7.6 and 8.7 we argued that the implementation was a satisfactory representation of the reduction, and thus that simulation results demonstrated our essential rules in action; by comparing the behaviour generated by the rules to the behaviour described in the specification we showed the reduction was sufficient to underlie the specification. Armed with essential rules that can underlie a real crowd behaviour, we are then in a position (as further described in section 4.3.2) to participate in a substantial way with other disciplines in refining our understanding of the behaviour of real crowds.

When we specify a reduction of the specification, it is natural that interesting factors fall by the wayside due to the abstraction process. For example, in chapter 8, the richness of cues to regulate front-to-back communication was dropped in favour of a single force-based mechanism. We found that the formal statement of the reduction highlighted those interesting pieces of future work. In a sense, the reduction process identified those parts of the behaviour that were of most importance and those parts of the behaviour which were not essential but were of interest, and separated these from the details that were of less interest. This was a benefit to using the methodology.

9.2.3 Effect on implementation

Because the reduction level must be specified from the point of view of an individual within the crowd, it supports implementations that are pre-disposed to take the perspective of individuals in determining actions (like movement decisions). This encourages a focus on situated processing and the modelling of behaviour that is determined based on information available to an individual at a specific time. It also guides the selection of an appropriate implementation formalism. For example, a mathematical model that deals with agents in aggregate is not appropriate to the individual and situated focus of a microscopic crowd simulation. When crowd phenomena

are expressed in the language of individual behaviours the non-interchangeability of agents becomes apparent and aggregate models simply appear infelicitous. This is the reason to argue for the appropriateness of the formalism as a part of validating the model.

In the case of the demonstrations of chapters 6–8, this important function of the methodology guided us in restating the floor field model as a multi-agent system. This provided a highly expressive implementation formalism that was well-suited to modelling crowd behaviour from the perspective of the individual. The appropriateness of the formalism was evident in all three demonstrations as the rules provided by the reduction were effectively added right at the heart of the floor field model. For example, in chapter 7, our changes to the model were tightly integrated with the core equation of the model by aligning the implementation of the reduction with the structures of the existing model.

It should be noted that in our demonstrations we were able to build on the results of previous chapters. In chapter 7 we were able to build on the results of chapter 6. In chapter 8 we were able to build on the results of chapter 7. This ability to use the methodology in a serial manner to progressively extend a model with multiple human factors is not a guaranteed property of all models developed using the methodology. In other words, the methodology does not guarantee that resultant models will be suitable for further expansion, even if the goal is to add multiple behaviours that are known to co-exist in crowds. This is due to the principles of multiple realisability discussed in section 4.3.2. Thus, each time the methodology is used to generate an additional behaviour, the appropriateness of the base formalism must be re-assessed. In other words, it is important to determine, each time a new behaviour is added to a microscopic model, whether that model is most appropriate implementation formalism — even if that model has itself been created according to the principles of the methodology.

9.3 Sime's concerns

The motivation for the Microscopic Human Factors methodology was ultimately to take account of Sime's concerns within the domain of microscopic modelling. Let us repeat Sime's quotation

from page 4 from which we began the work of this thesis. Sime was motivated by a lack of balance in research of crowd behaviour, saying:

Psychology and engineering are characteristically mutually exclusive in their focus on the perspective of crowd members who *think* and *behave* (psychology) or on static and dynamic objects (engineering) . . . The history of psychology, and indeed crowd psychology, is one in which the physical environment in which behaviour takes place has been ignored. While this might justifiably frustrate engineers, it is equally the case that social scientists are bemused by the lack of attention in engineering to the psychological and social aspects of crowd safety. . . (1995)

We have developed the Microscopic Human Factors methodology for the purpose of introducing human behaviour into microscopic crowd modelling, which often has a statistical-mechanics focus (corresponding to engineering in Sime's comment). As just described, the methodology is specifically intended to develop human-focused models (corresponding to psychology in Sime's comment), by starting with a description of how a behaviour works in real crowds. The cognitive focus continues as the methodology requires a complete description of the behaviour from the perspective of an individual. By requiring a close integration with a microscopic model, the methodology provides for the introduction of human factors at the heart of a simulation. This yields a simulation that is naturally situated and takes account of Sime's concerns; by maintaining the microscopic focus on spatial realism while introducing human behaviour, the methodology helps to unify human and engineering approaches to crowd modelling.

As discussed in section 1.1.2, panic is a problematic concept in crowd dynamics. The level of the specification in the methodology requires a clear statement of the researcher's concept of the crowd behaviour. This specification must not contain elements of panic, a responsibility that must be undertaken by modellers. By focusing on individuals rather than on collective structures at the reduction level, however, models of panic are also discouraged by the methodology.

Each of chapters 6–8, the demonstrations of the methodology, also provide a demonstration of Sime's concerns being addressed. In chapter 6, for example, we introduced human behaviour into the model in the shape of voluntary pushing behaviours; these behaviours did not emerge mechanistically as a general consequence of masses of people, rather they were rooted in a hypothesis about individual behaviour and frustrated movement. In chapter 7 we introduced

internal representations of the world, physical discovery of new spatial facts and communication of information; this chapter implemented Sime's suggestion that people should be viewed as moving through an information system as well as a physical environment. In chapter 8 we introduced communication behaviours in which people communicated on an individual basis to reduce dangerous crowd forces, resulting in global reduction in injury. In all cases, using the methodology, we were able to capture essentials of the human behaviour, and introduce it into the heart of the microscopic model.

Each of the demonstrations has: shown the value of the methodology in generating new insights into crowd behaviour, demonstrated the practicality of the model for generating microscopic representations of human behaviour, shown how to clarify the intellectual position taken by a model with respect to these behaviours, and improved the model with respect to Sime's concerns. We conclude that the methodology is a practical and useful way to improve microscopic crowd models.

Chapter 10

Conclusions and future work

Jonathan Sime expressed his concerns that behavioural and physical research into crowd dynamics were not adequately integrated. Although years have passed since his statements, and models have become more capable, his core concerns have not been well addressed. In this thesis, we have presented a methodology for introducing behavioural factors into microscopic models of crowds. The Microscopic Human Factors methodology helps to identify a human behaviour, to reduce it to its essential core, and to implement it at the heart of a microscopic model. We have presented three demonstrations of this process using the methodology, and derived conclusions and suggestions for future work from each demonstration. In general discussion we have combined results from the demonstrations with the literature to support conclusions about crowd safety, before drawing general conclusions about the usefulness of the methodology and its applicability to Sime's concerns.

In this chapter we present the conclusions obtained from the preceding work, we summarise the contributions made by this thesis, and we identify the future work it suggests. Finally, we shall close with a few concluding remarks.

10.1 Conclusions

The following conclusions have been generated by the previous chapters of this thesis:

1. Current microscopic models do not take account of Jonathan Sime's concerns relating to the integration of human behaviour and physical approaches to understanding crowd dynamics.
2. Using our Microscopic Human Factors methodology it is possible and practical to build microscopic models that can take account of Sime's concerns, generate interesting qualitative discussion and useful conclusions.
3. The floor field model (Burstedde et al., 2001; Kirchner & Schadschneider, 2002), reformulated as a multi-agent system, is an excellent match for the Microscopic Human Factor methodology and was easily fused with microscopic human behaviour simulations.
4. Small numbers of agents at the rear of crowds can influence agents at the front by triggering and focusing pushing and leaning forces, creating injuries and jamming exits.
5. Smooth reductions in numbers of injuries belie complexity in modelled exit rates due to a variety of protective injury effects.
6. As a consequence of conclusions 4 and 5, it is essential to model force in crowd models.
7. Our results support the view that medium drive produces the best exit rates in most cases. This accords with Proulx and Sime's work regarding the desirability of medium drive during fire egress (1991).
8. Differing agent knowledge produces a heterogeneous crowd, and is effectively engendered through local information discovery.
9. In a heterogeneous crowd, situatedness remains important, producing differences in behaviour that are as great for agents with similar knowledge as they are for agents with different information.
10. Modelling crowd heterogeneity through information discovery and communication increases the range of scenarios that can be modelled microscopically, including complex scenarios involving beliefs that change with time.

11. Heterogeneous movement at high drive can create obstructive barrier formations reminiscent of the one seen at the Lenin Stadium disaster.
12. Communication of spatial information and consequent reduction of differences between groups has the potential to be a powerful force mediator and enhancer of crowd safety.
13. Front-to-back communication can in principle reduce damaging forces, as demonstrated in a laboratory scenario, in our reconstructed Riverfront Coliseum disaster and in barrier formations.
14. Our results suggest that injuries may have occurred at Riverfront Coliseum when two perpendicular forces produced a hot spot of pressure. This differs from Fruin's suggestion (1993) that injury-causing forces were opposed to each other.

10.2 Contributions

The work in this thesis makes the following original contributions to knowledge:

1. We have identified microscopic behaviour simulations as a way to take account of Sime's concerns (that psychology and engineering are not adequately integrated in the study of crowd dynamics) in the domain of microscopic crowd modelling.
2. We have created the Microscopic Human Factors methodology to enable microscopic modellers to take account of Sime's concerns, integrating human behaviour and physical realism without losing the benefits of current microscopic models.
3. We have established a method to argue for the validity of models built according to the methodology, its benefits and its limitations.
4. We have contributed a force simulation to the floor field model (Burstedde et al., 2001) and its modelling community, allowing the modelling of injuries, crowd safety and force's effects on movement.

5. We have contributed a simulation of spatial information discovery and communication to the floor field model and its modelling community, providing for richer scenarios and crowd heterogeneity through creation of groups formed by shared mental content, as well as situatedness.
6. We have provided a conception of — and proposed a simple mechanism for — successful front-to-back communication.
7. We have contributed a simulation of front-to-back communication to the floor field model and its modelling community.
8. We have demonstrated that successful front-to-back communication can manage forces and reduce injuries in a laboratory model, in a reconstruction of the Riverfront Coliseum disaster and in a scenario similar to the barrier formation that caused deaths at Lenin Stadium.

10.3 Future work

The work of chapters 6–9 has raised many interesting points for future work. This section summarises them briefly, and the reader is referred to the full discussion for details.

1. We have suggested that the floor field model, when augmented by force, could benefit from diagonal movement in order to resolve conflicts at high agent drive.
2. Enhancing the force simulation with different classes of injury could improve the model. In real crowds injuries can involve slower movement (rather than no movement) and can present differing amounts of obstruction to other agents.
3. One question not answered in the literature at present concerns why people initiate voluntary pushing in the first place, a more nuanced simulation of the trigger for pushing would be of great value in understanding why some crowd situations are more physical than others.

4. We have suggested that injuries act as force breaks within the crowd. It would be better to install physical force barriers than to injure people. Accordingly, studying the most effective barrier configurations, and the effectiveness of fixed configurations to regulate force patterns generated by different densities and drive levels would be of great interest.
5. We have noted the presence of falls and seed injuries in the explanations of the Bethnal Green and Riverfront Coliseum disasters. It would be interesting to investigate this by creating seed injuries and evaluating their effects on other agents.
6. We have explored communication of spatial information under idealised circumstances in which agents communicate only when blocked in movement. There are many different reasons why agents might choose to communicate (e.g. altruism, leadership) and it would be interesting to experiment with these as well. In addition, local or global overhead announcements could further alter the pattern of behaviour.
7. In spatial communication introduction of probabilities for communicating and believing incoming communication could allow determination of the degree of compliance required to see the benefits of communication. This would allow determination of whether these benefits degrade gracefully or have a threshold below which they are no longer seen.
8. The ICE approach to building safety (Gwynne & Boswell, in press; Gwynne & Kuligowski, in press) discusses the importance of understanding what people learn about structures during normal usage in order to understand egress behaviours. An abstract simulation of this phenomenon could be implemented using the work of chapter 7 as a base.
9. Front-to-back communication has much in common with spatial communication, and would benefit from similar future work. In the case of front-to-back communication adding other cues to initiate the behaviour would be of interest, as would be the regulation of initiation, honouring and retransmission through probabilities, interaction with local and global overhead announcements and looking at the effects of pockets of successful and unsuccessful communication.

10. We have raised the question of training crowds to develop effective front-to-back communication, and suggested this could be of particular interest in large, slow situations that might also provide an opportunity for instruction. Determining the feasibility and practicality of this approach is left to future work.
11. Issues of timing in front-to-back communication seem extremely relevant to its utility in limiting injuries. This suggests separating the time progression of force, communication and movement (which currently are synchronised) to determine any timing related effects.
12. We have noted the potential for transitions from low \rightarrow high drive to be a better predictor of injury than simple failure of front-to-back communication. This suggestion could be examined by allowing dynamic variation of the floor field model's k_S and k_D parameters with time or discovery.
13. Building on suggestions 11 and 12, we have noted that two of our disaster scenarios featured environmental cues that could have synchronised transitions between low and high drive. It would be interesting to review a number of crowd disasters looking for this type of synchronised transition. This also suggests adding a study of synchronised transitions to the work of suggestion 12.
14. We have suggested that developing models of pre-movement time (a daunting but high-impact task) be considered a priority. It remains a key area of Sime's concerns, and we are not aware of crowd models that have successfully addressed it.

10.4 Closing remarks

In section 4.1 we presented our research objective:

Research objective: *To develop a methodology to integrate relevant human behaviours with physical factors at the heart of microscopic crowd models for the purpose of improving relevance to people and studying emergent crowd-level effects.*

The points in this chapter give credence to our success in meeting the objective. We have developed a novel methodology for integrating microscopic human factors with the existing physical simulation in microscopic models. We have provided three demonstrations of this methodology in use in the floor field model; each has been novel and generative in its own right in results, discussion and future work. We have argued that the methodology supports the addition of human behaviours to microscopic models, improving the relevance of these models to people. We have improved the models according to Sime's concerns, and also by increasing the complexity of the situations they can model.

Merton long ago described his attraction to what he called middle-range theories in social science (1949), a view later distilled as a challenge “neither to produce big, broad and general theories of everything, nor to spend time in empirical accounts *per se*, but to formalise, test, use and extend theoretical models able to shed light on the causal mechanisms that are behind the complexity of empirical phenomena” (Boero & Squazzoni, 2005). This challenge accords perfectly with the task of microscopic models of crowd dynamics. The fact that it originated with a social scientist reminds us that as modellers — despite the physical factors inherent in all crowd situations (safe or dangerous) — the complexity of empirical phenomena in crowd dynamics must be understood in the context of the human behaviour driving the individuals who make up the crowd. In other words: crowds are made of people.

References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- Averill, J. D., Mileti, D., Peacock, R., Kuligowski, E. D., Groner, N., Proulx, G., et al. (2005). *Occupant behavior, egress and emergency communications*. Washington: National institute of standards and technology.
- Axelrod, R. M. (1984). *The evolution of cooperation*. New York: Basic.
- Bandini, S., Federici, M., Manzoni, S., & Vizzari, G. (2006). Towards a methodology for situated cellular agent based crowd simulations. In O. Dikenelli, M.-P. Gleizes, & A. Ricci (Eds.), *Engineering societies in the agents world VI* (p. 203-220). New York: Springer.
- Bandini, S., Manzoni, S., & Simone, C. (2002). Heterogeneous agents situated in heterogeneous spaces. *Applied artificial intelligence*, 16, 831–852.
- Batty, M., Desyllas, J., & Duxbury, E. (2003). Safety in numbers? modelling crowds and designing control for the Notting Hill carnival. *Urban Studies*, 40, 1573-1590.
- BBC. (2005). *Iraq stampede deaths near 1,000*. August 31, 2005. Retrieved 2007-09-20 from http://news.bbc.co.uk/2/hi/middle_east/4199618.stm.
- Blue, V., & Adler, J. (1998). Emergent fundamental pedestrian flows from cellular automata microsimulation. *Transportation research record*, 1644, 29-36.
- Blue, V., & Adler, J. (2000). Modeling four-directional pedestrian movements. *Transportation research record*, 1710, 20-27.
- Blue, V., & Adler, J. (2001). Cellular automata microsimulation for modeling bi-directional pedestrian walkways. *Transportation research record B*, 35, 293–312.
- Boero, R., & Squazzoni, F. (2005). Does empirical embeddedness matter? Methodological issues on agent-based models for analytical social science. *Journal of Artificial Societies and Social Simulation*, 8(4), 6.
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1991). *Swarm intelligence: From natural to artificial systems*. New York: Oxford University Press.

- Burstedde, C., Kirchner, A., Klauck, K., Schadschneider, A., & Zittartz, J. (2002). Cellular automaton approach to pedestrian dynamics – Applications. In M. Schreckenberg & S. Sharma (Eds.), *Pedestrian and evacuation dynamics* (pp. 87–98). New York: Springer.
- Burstedde, C., Klauck, K., Schadschneider, A., & Zittartz, J. (2001). Simulation of pedestrian dynamics using a two-dimensional cellular automaton. *Physica A*, 295, 507-525.
- Chertkoff, J., & Kushigian, R. (2001). *Don't panic: The psychology of emergency egress and ingress*. London: Praeger.
- Cincinnati Enquirer. (1999). Photo of the lobby at Riverfront Coliseum with bodies and attendants in the foreground. Originally published December, 1979. Retrieved 2008-01-31 from http://www.enquirer.com/editions/1999/11/28/who_600x442.jpg. Re-publication was with article Nager, L. N. (1999, Dec. 3). Concert industry learned from Who tragedy, *Cincinnati Enquirer*, retrieved 2008-01-31 from http://www.enquirer.com/editions/1999/12/03/loc_concert_industry.html.
- Clarke, L. (2001). Panic: Myth or reality? *Contexts*, 1, 21–26.
- Crowd Dynamics. (2007). *Crowd disasters*. Retrieved 2008-08-10 from <http://www.crowddynamics.com/Main/Crowddisasters.html>.
- Dawson, M. R. W. (1998). *Understanding cognitive science*. Oxford: Blackwell.
- Dunne, L. R. (1999). *Tragedy at Bethnal Green 1943*. London: The stationery office. (Reprinted from *Tragedy at Bethnal Green: Report on an inquiry in to the accident at Bethnal Green tube station shelter*, 1943)
- Edelman, R. (1993). *Serious fun: A history of spectator sports in the USSR*. New York: Oxford.
- Fahy, R. (1994). EXIT89 — an evacuation model for high-rise buildings: Model description and example applications. In T. Kashiwagi (Ed.), *Proceedings of the 4th international symposium: international association for fire safety science* (p. 657-668).
- Feinberg, W. E., & Johnson, N. R. (2001). Primary group size and fatality risk in a fire disaster. In *Human behaviour in fire (2nd international symposium)* (pp. 11–22). London: Interscience.
- Festinger, L. (1954). A theory of social comparison processes. *Human relations*, 117–140. Cited in (Fridman & Kaminka, 2006).
- Fischer, P., Greitemeyer, T., Pollozek, F., & Frey, D. (2006). The unresponsive bystander: Are bystanders more responsive in dangerous emergencies. *European journal of social psychology*, 36, 267–278.
- Fridman, N., & Kaminka, G. A. (2006). Modelling crowd behavior based on social comparison theory. In S. El Yacoubi, B. Chopard, & S. Bandini (Eds.), *Cellular automata: 7th international conference on cellular automata for research and industry (ACRI 2006)* (pp.

- 694–698). New York: Springer.
- Fridman, N., & Kaminka, G. A. (2007). Towards a cognitive model of crowd behavior based on social comparison theory. In *Proceedings of the twenty-second national conference on artificial intelligence (AAAI-07)* (pp. 731–737). Menlo Park, CA: AAAI Press.
- Fruin, J. J. (1971). *Pedestrian design and planning*. New York: Metropolitan association of urban designers and environmental planners.
- Fruin, J. J. (1993). The causes and prevention of crowd disasters. In R. A. Smith & J. F. Dickie (Eds.), *Engineering for crowd safety*. New York: Elsevier.
- Gopnik, A. (1999). Theory of mind. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences* (pp. 838–840). Cambridge, MA: MIT Press.
- Gwynne, S. M. V., & Boswell, D. (in press). The use of a structure and its influence on evacuation behavior: FIRE & ICE. In *Pedestrian and evacuation dynamics 2008*.
- Gwynne, S. M. V., & Kuligowski, E. D. (in press). Application modes of egress simulation. In *Pedestrian and evacuation dynamics 2008*.
- Hardin, G. (1968). The tragedy of the commons. *Science*, *162*, 1243–1248.
- Helbing, D., Farkas, I. J., Molnár, P., & Vicsek, T. (2002). Simulation of pedestrian crowds in normal and evacuation situations. In M. Schreckenberg & S. Sharma (Eds.), *Pedestrian and evacuation dynamics* (pp. 21–58). New York: Springer.
- Helbing, D., Farkas, I. J., & Vicsek, T. (2000a). Freezing by heating in a driven mesoscopic system. *Physical review letters*, *84*, 1240–1243.
- Helbing, D., Farkas, I. J., & Vicsek, T. (2000b). Simulating dynamical features of escape panic. *Nature*, *407*, 487–490.
- Helbing, D., Johansson, A., & Al-Abideen, H. Z. (2007). The dynamics of crowd disasters: An empirical study. *Physical review E*, *75*, 046109.
- Henein, C. M., & White, T. (2005). Agent-based modelling of forces in crowds. In P. Davidsson, B. Logan, & K. Takadama (Eds.), *Multi-agent and multi-agent-based simulation* (pp. 173–184). New York: Springer.
- Henein, C. M., & White, T. (2006). Information in crowds: The swarm information model. In S. El Yacoubi, B. Chopard, & S. Bandini (Eds.), *Cellular automata: 7th international conference on cellular automata for research and industry (ACRI 2006)* (pp. 703–706). New York: Springer.
- Henein, C. M., & White, T. (2007). Macroscopic effects of microscopic forces between agents in crowd models. *Physica A*, *373*, 694–712.

- Henein, C. M., & White, T. (in press-a). Front-to-back communication in a microscopic crowd model. In *Pedestrian and evacuation dynamics 2008*.
- Henein, C. M., & White, T. (in press-b). The microscopic model and the panicking ball-bearing. In *Pedestrian and evacuation dynamics 2008*.
- Johnson, N. R. (1987a). Panic and the breakdown of social order: popular myth, social theory, empirical evidence. *Sociological focus*, 20, 171–183.
- Johnson, N. R. (1987b). Panic at the "Who Concert Stampede": An empirical assessment. *Social Problems*, 34, 362–373.
- Keating, P. J. (1982). The myth of panic. *Fire Journal*, 76(3), 57-61.
- Kirchner, A., Kluepfel, H., Nishinari, K., Schadschneider, A., & Schreckenberg, M. (2003). Simulation of competitive egress behavior: Comparison with aircraft evacuation data. *Physica A*, 324, 689–697.
- Kirchner, A., Nishinari, K., & Schadschneider, A. (2003). Friction effects and clogging in a cellular automaton model for pedestrian dynamics. *Physical review E*, 67, 056122.
- Kirchner, A., & Schadschneider, A. (2002). Simulation of evacuation processes using a bionics-inspired cellular automaton model for pedestrian dynamics. *Physica A*, 312, 260-276.
- KLM. (2007). *KLM main sponsor benefit match Spartak Moscow - HFC Haarlem*. Press release from KLM. Retrieved 2007-09-04 from <http://www.sportweek.nl/spartak-haarlem-english-klm.asp>.
- Kretz, T. (2007). *Pedestrian traffic*. Unpublished doctoral dissertation, Duisberg-Essen University.
- Kretz, T., Boenisch, C., & Vortisch, P. (in press). Comparison of various methods for the calculation of the distance potential field. In *Pedestrian and evacuation dynamics 2008*.
- Kretz, T., & Schreckenberg, M. (2006). The F.A.S.T.-Model. In S. El Yacoubi, B. Chopard, & S. Bandini (Eds.), *Cellular automata: 7th international conference on cellular automata for research and industry (ACRI 2006)* (pp. 712–715). New York: Springer.
- Lack, J. (2003, Feb 15). The hush-hush catastrophe. *The Guardian*. Retrieved 2007-09-03 from <http://www.guardian.co.uk/weekend/story/0,3605,894649,00.html>.
- Latané, B., & Nida, S. (1981). Ten years of research on group size and helping. *Psychological Bulletin*, 89, 308–324.
- LeBon, G. (1960). *The crowd*. New York: Viking. (Original work published 1895)
- Mamei, M., Zambonelli, F., & Leonardi, L. (2002). Co-Fields: Towards a unifying approach to the engineering of swarm intelligent systems. In P. Petta, R. Tolksdorf, & F. Zambonelli (Eds.), *Engineering societies in the agents world III, third international workshop, ESAW*

- 2002 (pp. 68–81). New York: Springer.
- Mamei, M., Zambonelli, F., & Leonardi, L. (2004). Co-Fields: A physically inspired approach to motion coordination. *Pervasive computing*, 3, 52–61.
- Merton, R. K. (1949). *Social theory and social structure*. New York: Free press.
- Musse, S. R., & Thalmann, D. (2001). Hierarchical model for real-time simulation of human crowds. *IEEE transactions of visualization and computer graphics*, 7, 152–164.
- Nguyen, Q.-A. H., McKenzie, F. D., & Petty, M. D. (2005). Crowd behaviour cognitive model architecture design. In *Proceedings of the 14th conference on behaviour representation in modeling and simulation (BRIMS-05)*.
- Nishinari, K. (in press). Towards smooth movements of crowds. In *Pedestrian and evacuation dynamics 2008*.
- Nishinari, K., Kirchner, A., Namazi, A., & Schadschneider, A. (2004). Extended floor field CA model for evacuation dynamics. *IEICE Transactions on information and systems*, E87-D(3), 726-732.
- Nishinari, K., Kirchner, A., Namazi, A., & Schadschneider, A. (2006). Modelling of self-driven particles: foraging ants and pedestrians. *Physica A*, 372, 132–141.
- North, M., Collier, N. T., & Vos, J. R. (2006). Experiences creating three implementations of the repast agent modeling toolkit. *ACM Transactions on modeling and computer simulation*, 16(1), 1–25.
- Oxford english dictionary* (2nd ed.). (1989).
- Oxford english dictionary* (Mar. 2007 draft ed.). (2007).
- Pan, X. (2006). *Computational modeling of human and social behaviors for emergency egress analysis*. Unpublished doctoral dissertation, Stanford University.
- Pan, X., Han, C., & Law, K. (2005). A multi-agent based simulation framework for the study of human and social behavior in egress analysis. In *Proceedings of the ASCE 2005 international conference on computing in civil engineering*. Retrieved 2007-06-20 from http://eil.stanford.edu/xpan/ASCE_2005_Pan_Han_Law.pdf.
- Pauls, J. (1984). The movement of people in buildings and design solutions for means of egress. *Fire Technology*, 20, 27-47.
- Pauls, J., Fruin, J. J., & Zupan, J. M. (2007). Minimum stair width for evacuation, overtaking movement and counterflow. In N. Waldau, P. Gatterman, H. Knoflacher, & M. Schreckenberg (Eds.), *Pedestrian and evacuation dynamics 2005* (pp. 57–69). New York: Springer.

- Pelechano, N., Allbeck, J. M., & Badler, N. (2007). Controlling individual agents in high-density crowd simulation. In *Eurographics / ACM SIGGRAPH symposium on computer animation*. New York: ACM.
- Pelechano, N., O'Brien, K., Silverman, B., & Badler, N. (2005). Crowd simulation incorporating agent psychological models, roles and communication. In *V-crowds 2005*. Retrieved 2006-04-26 from http://www.seas.upenn.edu/~npelecha/Pelechano_V_Crowds05.pdf.
- Proulx, G. (2002, Fall). Cool under fire. *Fire protection engineering* (16), 23–25.
- Proulx, G., & Sime, J. D. (1991). To prevent 'panic' in an underground emergency: why not tell people the truth? In G. Cox & B. Langford (Eds.), *Fire safety science: 3rd international symposium* (pp. 843–852). London: Elsevier Applied Science.
- Quarantelli, E., & Dynes, R. R. (1972, Feb). When disaster strikes. *Psychology today*, 66–70.
- Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. In M. C. Stone (Ed.), *Proceedings of the 14th annual conference on computer graphics and interactive techniques* (pp. 25–34). New York: ACM.
- [Riverfront Coliseum] (n.d.) Unattributed aerial photo of the plaza of Riverfront Coliseum. Retrieved 2008-01-31 from <http://www.biwa.ne.jp/~presley/concert/river.jpg>.
- Ruizenaar, T. (2007, Feb 22). Spartak and haarlem mark anniversary of tragedy. *Reuters news*. Retrieved 2007-09-04 from <http://www.reuters.com/articlePrint?articleId=USL21339792200070223>.
- Schellinck, J. A. (2008). *A general perception based framework for modelling animal aggregation*. Unpublished doctoral dissertation, Carleton University, Ottawa.
- Sime, J. D. (1983). Affiliative behaviour during escape to building exits. *Journal of environmental psychology*, 3, 21–41.
- Sime, J. D. (1985). Designing for people or ball-bearings? *Design Studies*, 6, 163–168.
- Sime, J. D. (1995). Crowd psychology and engineering. *Safety Science*, 21, 1–14.
- Smith, B. C. (1999). Situatedness/embeddedness. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences* (pp. 769–770). Cambridge, MA: MIT Press.
- Still, G. K. (2000). *Crowd dynamics*. Unpublished doctoral dissertation, University of Warwick, U.K.
- Stillings, N., Feinstein, M. H., Garfield, J. L., Rissland, E. L., Rosenbaum, D. A., Weisler, S. E., et al. (1987). *Cognitive science: An introduction*. Cambridge: MIT Press.
- Wolff, M. (1973). Notes on the behavior of pedestrians. In A. Birenbaum & E. Sagarin (Eds.), *People in places: The sociology of the familiar* (pp. 35–48). New York: Praeger.