



Agent-based pedestrian modelling

Michael Batty

When the focus of interest in geographical systems is at the very fine scale, at the level of streets and buildings for example, movement becomes central to simulations of spatial activities. Recent advances in computing power and the acquisition of fine-scale digital data now mean that we can attempt to understand and predict such phenomena, with the focus in spatial modelling changing to dynamic simulations of the individual and collective behaviour of individual decision making at these scales.

In this chapter, we develop ideas about how such phenomena can be modelled showing first how randomness and geometry are all important to local movement and how ordered spatial structures emerge from such actions. We focus on these ideas with pedestrians, showing how random walks constrained by geometry but aided by what agents can see, determine how individuals respond to locational patterns. We illustrate these ideas with three examples: first for local scale street scenes where congestion and flocking is all important, second for coarser scale shopping centres such as malls where economic preference interferes much more with local geometry, and finally for semi-organised street festivals where management and control by police and related authorities is integral to the way crowds move.

Batty, M. (2003) Agent-Based Pedestrian Models, in P. Longley and M. Batty (Eds) *Advanced Spatial Analysis*, ESRI Press, Redlands, CA, 81-106.

1 Spatial dynamics at the fine scale: agents and infrastructure

As we approach scales where geometry in the form of streets, buildings and land parcels becomes significant, the focus of inquiry in urban modelling changes from one where location is dominant to situations where the local dynamics of movement are much more significant. This is easily illustrated in retailing. Firms fix their general location according to geodemographic patterns of demand and the availability and capacity of infrastructure, but their precise locations are determined according to the volumes of trade generated by their competitors. In turn, these depend not only on footfall in shopping streets but also on the ways in which goods are positioned relative to movement patterns of shoppers within the stores themselves. These ideas are manifest in the rules of thumb used by retail planners in organising the layout of shopping malls where stores position themselves according to the way their products relate to each other and to the position of the anchor stores that heavily influence the overall character of any given retail development.

Until quite recently these kinds of problem were seen as lying outside the remit of urban modelling and GIS. However, advances in computation—in terms of the acquisition of digital data using sensors, improved storage and processing of disaggregate data and advances in object-oriented programming that make it possible to represent and simulate large numbers of generic objects or agents—have enabled us to make substantial progress, some of which we discuss in this chapter. Conceptual developments dealing with dynamic systems at aggregate (see Batty and Shiode, this volume) and disaggregate levels (see Torrens, this volume) combined with ideas from complexity theory that enable us to link seemingly uncoordinated actions at the local level with the emergence of more global structures (see Longley et al, this volume) are providing the intellectual framework. Central to this is the idea of agents having behaviour that is simulated as a trade-off between selfish and unselfish actions, between individual and cooperative decision making. The classic case can be seen in local movement where the desire to reach some goal often leads to turbulence in crowd situations, while the desire to see what others are doing leads to flocking. Both can end either in congestion and panic or in smooth movement if local actions develop at speeds which allow synchronisation.

These models supplement rather than replace those used to model location and movement at coarser scales of granularity (see Longley et al, this volume) and add to the range of tools available to planners and engineers,

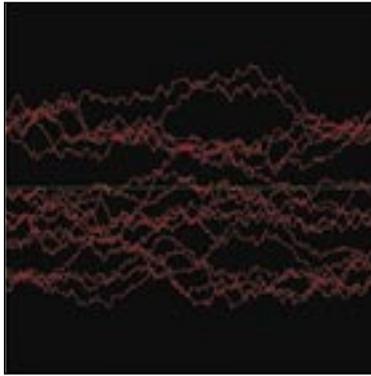
notwithstanding the need for strategies to enable their integration. They also pose important challenges to GIS for they involve geometry in a way that has not been central to spatial analysis, and they invoke a concern for dynamics that has long been an Achilles' heel of geography and GIS. The models that we will sketch out here all deal with pedestrian movement in streets, buildings and related complexes and introduce ideas about variability and heterogeneity in urban systems which are largely defined away in the more aggregative applications that characterise mainstream GIS applications. The way in which randomness, geometry, and economic and social preferences combine and collide will in fact be our starting point and this introduces a very different mix of modelling styles than any developed in this field hitherto. The nexus of this emerging field arises from the synthesis between randomness and geometry in physics, related to ideas concerning economic rationality and social preference.

We will begin with ideas about hypothetical walks—random walks, adding geometry and economy as we go, thus constructing a generic model of local movement. We illustrate the approach in the rest of the chapter by applying it to three very different situations: very fine-scale street situations dominated entirely by geometry and the collective phenomena of crowding and panic; more economically structured movements characteristic of shopping behaviour in centres and malls; and managed situations such as festivals and street parades where control is essential to public safety.

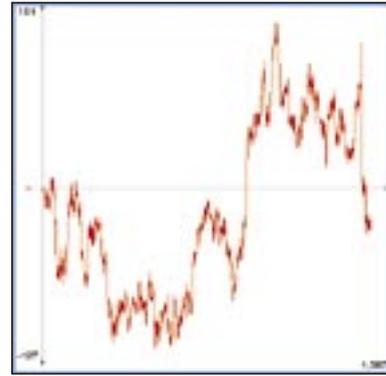
2 Random walks, geometry and locational preference

Before we define agent behaviour, we should say a little more about what agents actually are—for the ideas we will introduce in this section apply with equal force to aggregate as well as disaggregate behaviour. Agents in our definition are mobile and their behaviour relates largely to this mobility. This behaviour in human systems is not simply determined by preferences, intentions or desires. Rather, it is also determined by the environment that reflects the spatial or geometric structure in which the agents function as well as variability between agents, in terms of their intrinsic differences and the uncertainty that they have to deal with in making any response. This variability we treat as randomness. These could all apply to aggregate behaviour but at the disaggregate level, agent-based approaches have become popular in several fields. In the virtual world for example, there are software agents or 'bots'—pieces of code that are transmitted across networks and respond to requests which they are programmed to act upon; in the physical

world there are particles; in the natural world, plants and animals; and in our own domain—the human world—there are not only ourselves but also the agencies and institutions that we create to produce collective responses. The way in which randomness, geometry and intention combine is different in each of these domains. Modern physics, for example, is largely a product of randomness with the constraining effect of geometry, whereas modern economics is the product of randomness with distinct preferences that imply sentient behaviour. In human spatial systems at the level of the agent, behaviour is more complex in that it is clearly some product of randomness, geometry, economic intentions and social preference.



A The walk wraps around from left to right every 200 time periods



B The walk over its 1387 time periods

Figure 1 A one-dimensional random walk

We begin with the simplest one-dimensional random walk, in which an agent responds randomly to change in one dimension but moves consistently forward in the other. A typical example is a time series portraying the random variation of a stock or share price over time. Such a walk depends on the random variation from a preset position whose coordinates are calculated using the generic equation: new position = old (previous) position + random variation. In the case of a time series, the variation is in one direction through time where the coordinate of time is always known and increments inexorably while the change is random. Formally we can write $v(t+1) = v(t) + \epsilon_v$, where $v(t)$ is the value of the stock or related market index at time t and ϵ_v is the random error in the stock $v(t)$. This is a first order process where the lag in value of the stock is one time period (one order) and is often called a Markov process. In Figure 1 we show an example of the behaviour of this process where we have simply simulated variation from the starting point

of $v(0) = 0$ for some 1400 time periods. Every 200 time periods the value wraps around in Figure 1A, but we show the entire series in Figure 1B where it is clear that if you examine a small part of the series and then scale this up—aggregate it over time—the series is self-similar or fractal, a well-known characteristic of this kind of constrained randomness.

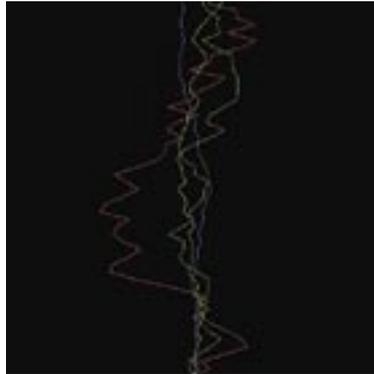
To think of this kind of walk as a walk in human space, we only need to think of Figure 1 as simulating how a human agent walks forward along the $y = 0$ line. Time thus becomes the forward direction and the randomness is a sideways deviation. If the apparent magnitude of the random variation in this figure suggests that our human agent is intoxicated, do remember that the amplitude of the random variation would appear much less if the y axis were marked in larger increments, and also that there is always some random deviation to help us steady ourselves when walking. Thus in Figure 2A, we show four different walks all from the same starting position where the blue walk has the least variation, the yellow a little more with the green more still: the red line might be associated with someone who has had too much to drink! In Figure 2B, we keep the variance the same but start a series of walkers from the same point and track their paths. Because change is random, some vary quite a lot, but there is a clear tendency for the walks to bunch together with the deviations from the straight line following a Normal distribution.

It is easy to write our generic model now for this kind of walk. Using x and y for the coordinates of each point on the walks in Figures 2A and 2B, then

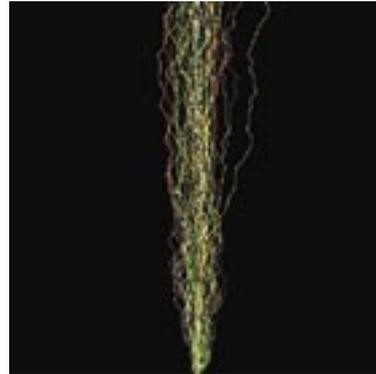
$$\left. \begin{aligned} x(t+1) &= x(t) + \varepsilon_x \\ y(t+1) &= y(t) + 1 \end{aligned} \right\} (1)$$

where it is very clear that our walk is highly constrained in the y direction. It is now simple to turn this walk into a two-dimensional random walk. In a sense it already is, but if we fix change in the y direction randomly just as we have done in the x direction, that is, we set $y(t+1) = y(t) + \varepsilon_y$, then we can simulate how the walkers walk all over the space not just in one direction across it. In Figure 2C we show the track of a walker who moves everywhere randomly but at least does not go outside the space—that is the walker avoids the edges—and in Figure 2D, we show the course of four walkers doing the same. If we let these walkers continue in this way, then at the scale of resolution at which we are viewing the space, all the pixels in the space will soon be visited and the 1-D tracks will come to fill the 2-D space. This is another example of randomness being fractal in that here is an

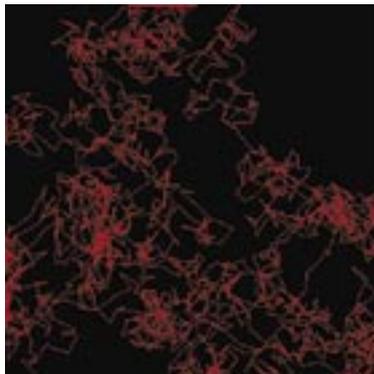
example of a space-filling curve—a line of one dimension filling a space of dimension two. We say that the resulting object has a Euclidean dimension of one but a fractal dimension of two.



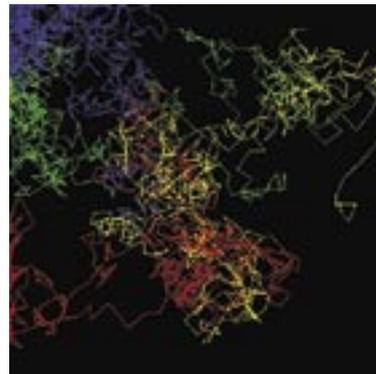
A 1-D walks—different random variances



B 1-D walks—same overall variances



C A 2-D edge-avoiding walks



D Four 2-D edge-avoiding walks

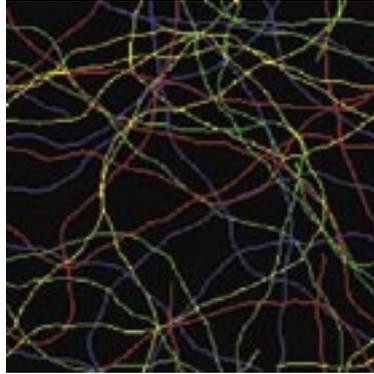
Figure 2 Multiple one- and two-dimensional random walks in the same 2-D space

This is not a chapter on geography and fractals, however, (see Batty and Longley 1994) and much as this may be of interest, it is not our central quest. We now need to show how the geometry of the environment begins to constrain and give structure to such random walks. First though, let us note that as soon as we begin to make assumptions about how the walker moves, we, unwittingly perhaps, introduce behaviour. The walks in Figure 2 twist too violently in terms of direction for these to be characteristic of the way human beings respond, and if we constrain such twists, producing the sort of gentler walks that we show in

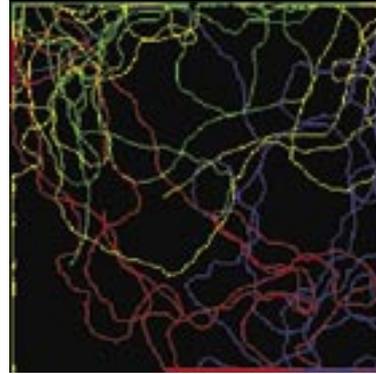
Figure 3A, then we simulate greater realism. It is clear that there are two very different ways in which the walk can be constrained to avoid crossing the edge of the space. The rather blunt way is that whenever a walker bumps against the edge, the walker then shifts away from the edge, while the more intelligent way is that the walker sees the edge long before it is reached and takes evasive action. The first way, shown in Figure 3B, simply prevents the walker from crossing the edge but is dominated by movement along the edge. The intelligent way requires a lot more computation. From any position the walker must look ahead and determine how far from the edge it is, then alter its heading in such a way that as it gets nearer to the edge, it veers away from it, thus retaining its smooth behaviour. We show such a strategy in Figure 3C where the grey lines are lines of sight computed from every point where the walker is located. Note how the track, shown in red, can be altered according to how far the walker is from the edge. This is akin to introducing vision into our system, but this is expensive in terms of computer time for it means that before a walker moves, it must have information about how far every possible move is from every possible obstacle.

Our last example involves the interaction of agents with one another. In Figure 3D we show how agents flock together from initial positions where they are entirely independent. Such flocking is based on the kind of curiosity involved in wondering what one's neighbour is doing, where s/he is going, and thence shadowing the neighbour's movement. Three factors determine how walkers flock. First, there is cohesion which is the direction pointing towards the centre of the neighbourhood containing the walkers who the walker in question is shadowing. Second, there is the separation in that a walker does not want to be too close to any of those in its neighbourhood and thus wishes to point away from those who are too close. And last, there is the alignment of walkers in their neighbourhood which is taken as the average direction in which all walkers are pointing. If these three criteria, which imply different headings, are weighted differentially different kinds of flocking take place and highly realistic simulations can be generated through their fine tuning. This kind of model was first proposed for 'boids', computer-simulated objects in flight, by Reynolds (1987) for simulating flocks of birds in computer movies, but has wide applicability as a model of the way groups of people coalesce and disperse in situations of crowding. We show the tracks of such behaviour for four agents who flock together and wrap from the top to the bottom of the screen in Figure 3D.

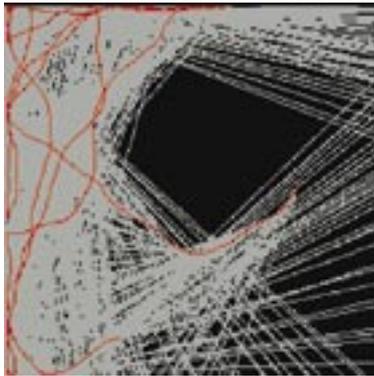
So far, we have shown how a random walk can be given more structure by responding to the geometry of its environment, but it is quite possible to create



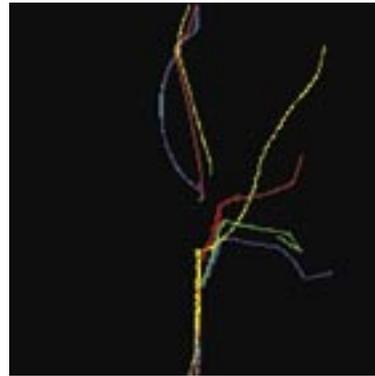
A Smooth paths—no edge avoidance



B Smooth paths—crude edge avoidance



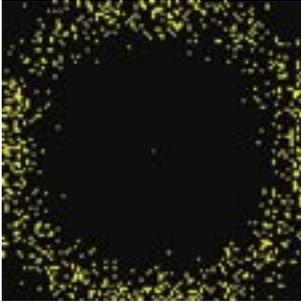
C Edge avoidance using vision



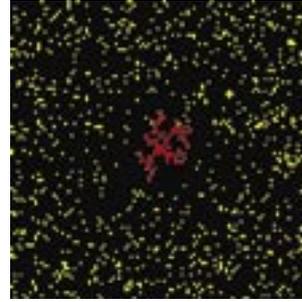
D Flocking with no edge avoidance

Figure 3 Smooth walks A, edge avoiding B, with vision C and with clustering-flocking D

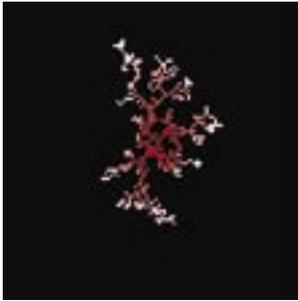
highly ordered geometric structures from the interaction of the random walkers themselves. In flocking, for example, the trail that is marked out becomes a dominant path from a set of unrelated paths in the first instance. There is, however, a dramatic example of this based on constraining each walk to terminate once the walkers reach any fixed point. For example, plant a fixed point at the centre of a space and then launch a series of walkers at a far distance from this point, letting them walk randomly in the space. As soon as one of the walkers ‘touches’ this point, it creates another adjacent point and then moves far away from the growing structure. If this process continues, what happens is that the structure grows in a treelike manner away from the initial point, creating a growing structure such as that shown in Figure 4 for three time periods.



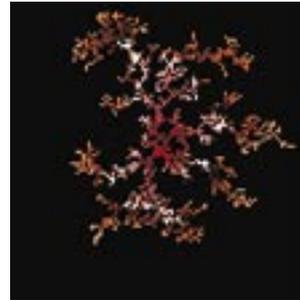
A The seed (red) at the centre; the agents(yellow) begin their random walk



B Creation of structure: when 100 agents have touched the growing mass



C The structure after 400 agents have come into contact with it



D The structure after 1000 agents have come into contact with it

Figure 4 The creation of structure: constrained diffusion-limited aggregation (DLA) from random walks

Basically, this is a good model of how a street system might grow in a town over a long time period. People migrate into the space and wander around until they come into contact with each other. Once they do so, they decide to settle and if the walkers are launched one at a time, then eventually the growing structure tends to fill the space available. The criterion for settlement is immediate adjacency to the existing structure and, as the walkers are more likely to come into contact with its edges (because the edges shield the centre of the structure), growth mainly takes place on the edge of the structure's growing tentacles. This is called diffusion-limited aggregation (DLA). At the most local level in terms of movement within streets, there is some sense in which such patterns reflect the location of shops on radial roads emanating from a town centre, thus reflecting location which is the product of movement (migration) over much longer time periods than those used in the models produced here. Once again such structures are fractals; this DLA

model being at the core of much recent work on thinking of city morphologies as fractals (Batty and Longley 1997).

3 Movement at the very fine scale: a generic model

We now need to stand back and derive a generic model for local movement from the ideas we have already introduced. All the elements for this have been stated, at least implicitly, and there are at least four features that direct movement. First, geometric obstacles need to be negotiated. Second, agents repel each other when congestion and crowding builds up, while third, agents are attracted to each—these second and third features being the elements used in flocking. The fourth element, though in one sense the most important, relates to the desired direction in which the walker wishes to travel. In our random walks, we assumed that this was either straight ahead or completely random, but in real situations, we must generate such directions as the product of preferences and intentions. A useful formulation of these ideas has been developed over several years by Helbing and his group (Helbing et al 2001). Helbing (1991) refers to his generic model of pedestrian movement as a social force model in which each of these four features is associated with a force that pushes the walker in a particular direction. In general, we might think of movement to a new location as being formed from

$$\begin{bmatrix} \text{new} \\ \text{position} \end{bmatrix} = \begin{bmatrix} \text{old} \\ \text{position} \end{bmatrix} + \begin{bmatrix} \text{desired} \\ \text{position} \end{bmatrix} + \begin{bmatrix} \text{geometric} \\ \text{repulsion} \end{bmatrix} + \begin{bmatrix} \text{social} \\ \text{repulsion} \end{bmatrix} + \begin{bmatrix} \text{social} \\ \text{attraction} \end{bmatrix} + \epsilon \quad (2)$$

where each of these components is some function that is ‘added’ into the algorithm that is used to compute movement. Desired position is a behavioural variable that relates to locational preferences while geometric repulsion is the summation of forces which stop a walker from bumping into some obstacle. Vision is important in its computation. Social repulsion and attraction are summations of all the interaction effects with other agents that are within the neighbourhoods or fields where such effects are relevant to movement.

There are several ways in which this formulation can be made operational. At very fine scales where panic and crowding are the main behaviours as in evacuation and emergency situations, the model can be formulated traditionally in terms of the physics of motion where velocity and acceleration play a central role (Helbing et al 2000). Where speed is not important then the coordinates of position can form the essence of computation (or headings if this be the preferred form: see

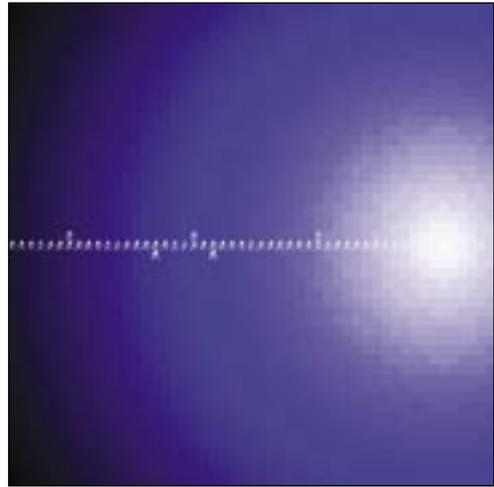
Batty et al 1998) while in situations where location is relative, then pixel location relative to local neighbourhoods can be used as in traditional spatial-interaction modelling (Batty et al 2003). In fact, in computation associated with such models, there is no simple equation structure, because motion must be computed through a sequence of decision rules that incrementally update position.

To show some simple examples of what this means, in Figure 5 we illustrate a narrow street where walkers are launched to the left (west) and stream towards a desired location at the right (east). Figure 5B shows that there is an attraction surface around the most desirable point, and one way of walkers being attracted is to compute their position according to the steepest gradients. The geometry of the street constrains their movement too and in Figures 5C and 5D, we show what happens when we reduce the variability—random movement—which is also an essential component of the way we handle crowding and dispersion. In Figure 6, we take this simple street and add obstacles by narrowing it. In Figure 6A, the street decreases in width immediately to a narrow outlet, while in Figure 6B and Figure 6C we show the effect of funnelling on the speed at which walkers are able to negotiate such a barrier. Finally, in Figure 6D, we show what happens if the street becomes in effect a series of rooms linked by narrow entrances and exits. This is the kind of analysis that these types of models are able to generate; Helbing et al (2000) have produced some extremely impressive results on the geometry of streets and corridors and the effects these have on different kinds of panic and its control.

In Figure 7, we show a situation where there are two types of pedestrian: those forming a parade which moves around the corner of a street intersection (agents in white) with pedestrians watching (agents in red). The model enables us to assess the build up in pressure between the two types of walker and to set a threshold which, if breached, leads to one crowd mixing with the other. This is achieved in Figure 7B where the red agents build up pressure in trying to see the parade (social attraction) but then try to diffuse (social repulsion) with serious consequences in that the only way this can happen is for the crowd to panic and break through the parade into freer space. This is the kind of disaster scenario that such models can be used to predict.



A Walkers move from left to right



B The attraction surface

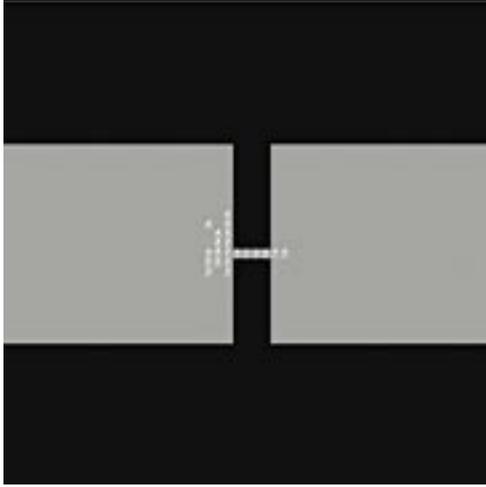


C Agents reaching the desired location

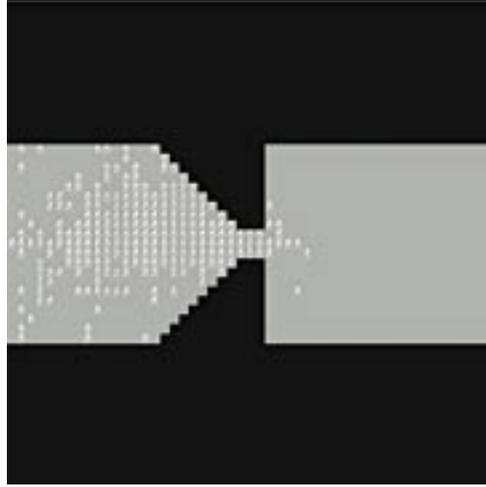


D Agents cluster at the desired location

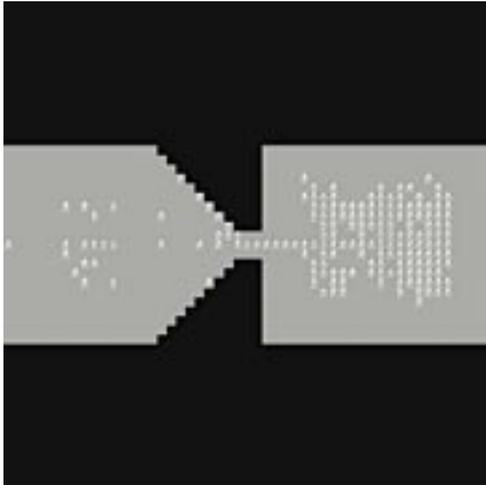
Figure 5 Walkers head towards the most desired location along a narrow street or corridor within limits of the centre line



A A narrow obstruction



B A funnelling obstruction



C Moving through the funnel



D A series of narrow obstructions

Figure 6 How obstructions delay movement, and compromise safety

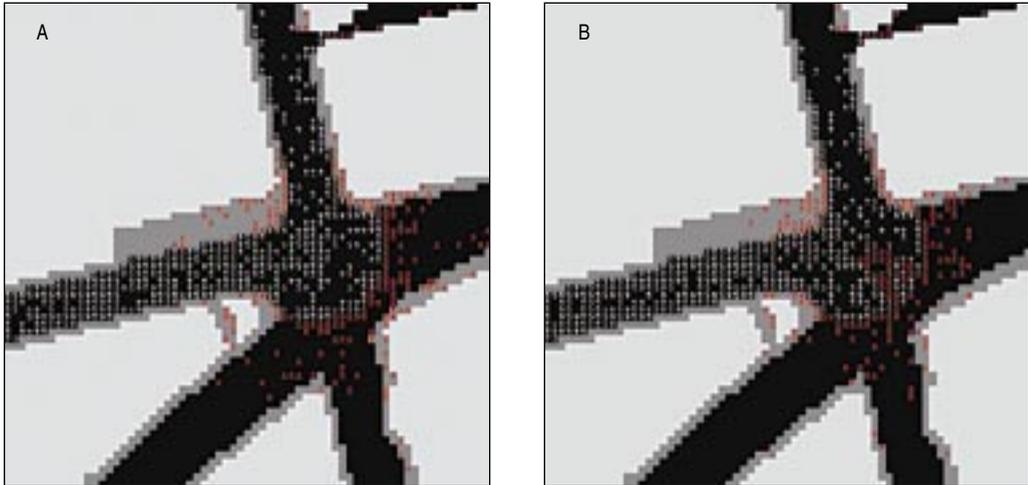


Figure 7 Panic as a crowd breaks through a cordon

The agents colored white are part of the parade while the red agents are those watching the parade. In A, these fringe the parade but in B, the pressure of the crowd on the paraders leads to them breaking into the parade.

4 Individual and collective behaviour: pedestrians in buildings, malls and centres

Our first application is to a complex building—the Tate Britain art gallery in London—that currently houses the gallery’s classical collection, although when our study was undertaken, the gallery also housed the modern collection now in the Tate Modern. The Space Syntax group at UCL (www.bartlett.ucl.ac.uk/research/space/overview.htm) specialises in recording pedestrian movement in buildings, and they carried out a survey of movement through the rooms of the gallery over a one-week period in August 1995. The building provides an excellent test bed for simulating local movement as there is only one entrance but many rooms and corridors which contain very different paintings and sculptures—thus providing a clearly differentiated set of destinations for art lovers and the more casual visitor alike. In many respects, the kind of movements associated with the gallery are not too dissimilar from those associated with browsing when shopping and thus the applications provided a rather good problem on which to fine tune the model.

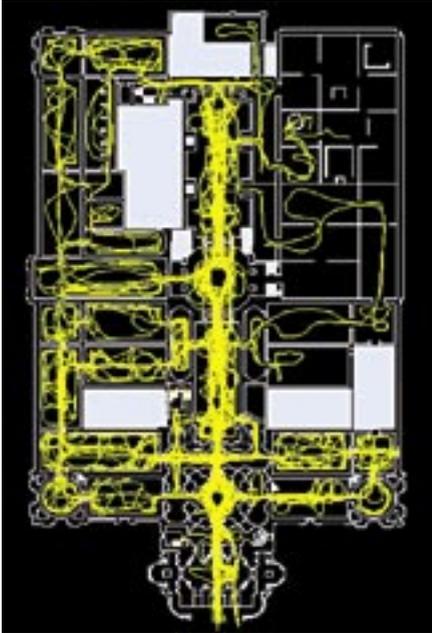
casual?

Figure 8A shows the pattern of movements recorded by the Space Syntax group over one hour in August 1995 using a technique of person following which the group had perfected in many other studies. It is clear from these patterns that the

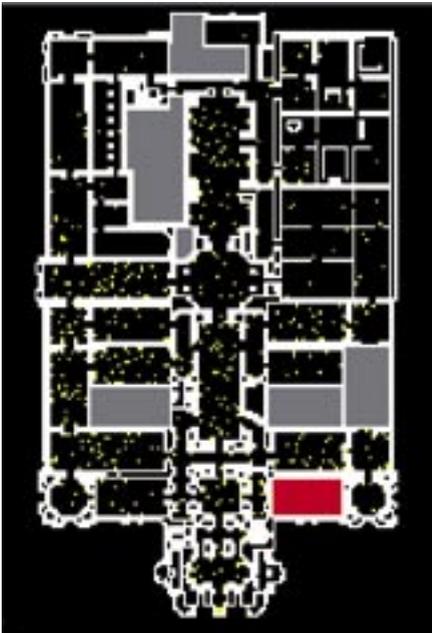
gallery has a greater number of visitors on the left side of its central axis where the classical and British collection was housed, in contrast to the rather sparser number of person movements associated with the right side where the modern collection was kept. It is unclear whether this reflects the preferences of the visitors or the more convoluted room structure associated with the right side—or indeed whether it reflects various perceptual clues associated with light in the gallery once they enter. The book shop—in red in Figure 8B—is also a major focus for movement having a somewhat different purpose from the rest of the gallery.

The model that we applied was built around Helbing's generic structure indicated in equation (2) on page 90 but with several variations to match the problem. To begin, the attraction surface was configured in two related parts: first, each room had a certain attraction while we also used a more global attractor to spread visitors throughout the gallery on the assumption that if a visitor found a room attractive, they would not spend all their time there but would wish to visit other areas as well. In terms of geometric repulsions, we developed two variants—one in which walkers moved right up to an obstacle and negotiated it in the rather blunt way by 'bumping' into it; and a second, much more elaborate version in which we incorporated vision into the structure by pre-processing. The pre-processing stage involved computing the visual (isovist) fields associated with each pixel point describing the gallery (Batty 2001). Social repulsion related directly to density limits on crowding and, if breached, resulted in dispersion within a local neighbourhood; social attraction was based on a flocking-like algorithm that enabled visitors to follow others if it appeared that flows to particular rooms were increasing. Fluctuations based on the random walk behaviour that we used in the 1-D walks in Figures 2A and 2B were also incorporated to reflect variations in walking behaviour within a large and heterogeneous population. We also experimented with some versions of the model where the number attracted to rooms changed the attraction surface associated with the rooms, starting with a version where all the rooms were equally attractive and seeing how a hierarchy of attraction evolved through movement. A typical output of the model is shown in Figure 8B where the yellow dots relate to the 550 or so walkers observed in the gallery during the one-hour period when the observations were made.

The model was operationalised using a series of decision rules, which was operated in strict sequence but within a parallel processing structure in which each agent made its own computations simultaneously with all the other agents. The essence of this model, however, is not to mimic the actual paths that were observed by the Space Syntax Group but to examine the statistical physics of steady state



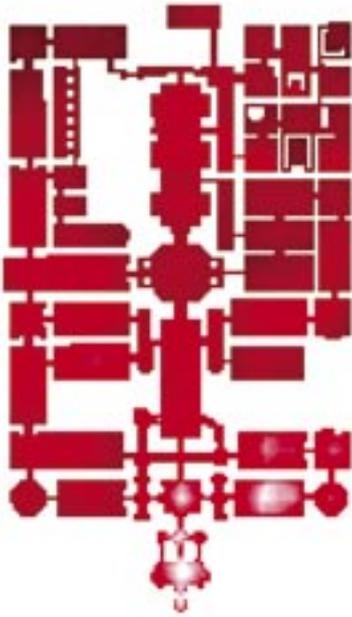
A Paths observed over one hour August 1995



B Simulating movement within and between rooms. (Red denotes bookstore).

Figure 8 Observation and simulation of visitors within the Tate Britain Gallery

In C, the density scale ranges from dark red at 1 persons per 20 square meters to white at 1 persons per 2 square meters



C Density of visits to rooms in the Gallery

behaviour associated with this kind of model in this kind of application. In Figure 8C we show the pattern of room densities which was computed over many runs of the model where the walkers moved around the gallery for many time periods. Here it is clear that from a starting position where all the rooms had equal attraction, the asymmetry of movement actually observed in the gallery was simulated by the model in its steady state. We can say little more than this, although the implication is that it is the configuration of the gallery rather than what is on its walls and in its rooms that conditions how people move within it. We can also use the model rather effectively to close and open certain rooms. The bookshop, for example, is a case in point (see the red room in Figure 8B), and if this is closed, then the pattern on the right-hand side of the gallery changes substantially.

The major problem that we face with any kind of pedestrian modelling is that the data that we are able to get are usually inadequate in many ways. Density counts are now easy from closed circuit television (CCTV), related sensors, or gate counts, but the observation of actual paths is fraught with difficulty. Along with colleagues at the Shibasaki Laboratory University of Tokyo, we are experimenting with laser scanning technologies, but although these are able to pick up actual paths travelled, there are limits on the density of the scene in terms of numbers of walkers. To illustrate that our models are much further ahead than

the data that we have, we have adapted the Tate model to a small town centre in the British Midlands — Wolverhampton — for which we have only crude footfall data. Nevertheless, this centre is rather simple as most of the movement within the centre is on foot with pedestrians coming from a ring of car parks around the edge of the centre as well as the rail and bus stations. There is little on-street parking and this makes the calibration of the model feasible. We show an example of the output from the model in Figure 9 where the streets within the town centre are clearly shown and where the movement originates from the car parks and stations around the edge. Wolverhampton is one of the few towns in England bounded by a complete ring road with no residents living in the centre, and thus all movement is associated with shopping and office work. The structure of the model differs from the Tate in that the global attraction surface is more complex albeit centred on the prime retail pitch. Local attractions are also built into the model based on the idea that once a shopper visits a particular shop, the probability of shopping in an adjacent shop is increased or decreased dependent upon that shop and the type of shopper. In short, movement, although conditioned by a general pattern of moving from edge to the centre, is complicated by the actual visits made. In this way the physical and locational structure of activities in the centre becomes associated with the preference of the shopper.

This framework has been elaborated in CASA with the well-known SWARM modelling system that was designed to simulate artificial life forms. Haklay et al (2001) developed such a model where preferences and geometry were combined in a much more sophisticated way than in the model shown here. In fact, in the model whose typical output is shown in Figure 9, shoppers only visit the centre for a fixed period of time with an increasing number returning to their rail and bus stations and car parks while others enter the centre. In this way, we are able to build the routine dynamics—the ebb and flow of shoppers during the shopping day—into the model and thus begin to assess critical densities at particular points in time. Thus we can begin to get a handle on locational attraction not only in space but also in time, and this leads us full circle to ideas about the 24-hour city and the way retailers are beginning to compete in time.

5 Highly managed spatial events: street parades and carnivals

To apply a full version of the model, we will now adapt it to a problem with major implications for crowd safety. Unlike shopping trips which have quite distinct purposes reflecting various preferences for places and goods, movements associated



Figure 9 Simulating pedestrian shopping trips from car parks in Wolverhampton Town Centre (within the ring road—mapped extent 1 km square)

with street festivals and parades are in a sense simpler, more focused locationally but in another sense, more complex in that walkers and visitors respond less predictably to the stimuli that together comprise the event. We have built a model of an annual street festival in west central London known as the Notting Hill Carnival. The Carnival has grown from a small West Indian street celebration first held in 1964 to a two-day international event attracting 710 000 visitors in 2001. It consists of a continuous parade along a circular route of nearly 5 kms in which 90 floats and 60 support vehicles move from noon until dusk each day. Within the 3 km² parade area, there are 40 static sound systems, and 250 street stalls selling food. The peak crowds occur on the second day between 4p.m. and 5p.m. when in 2001, there were some 260 000 carnival visitors in the area. In 2001 over the two days there were 500 accidents, 100 requiring hospital treatment with 30 per cent related to wounding, and 430 crimes committed with 130 arrests. Some 3500 police and stewards were required each day to manage the event.

The safety problems posed by the event are considerable. There are many routing conflicts because of cross movements between the parade and sound systems, while access to the carnival area from public transport is uneven with four roads into the area taking over 50 per cent of the traffic. A vehicle exclusion zone rings the area and thus all visitors walk to the carnival. Crowd densities are high, overall at about 0.25 persons per m²; a density of 0.47 ppm² line the carnival route, while there are 0.83 ppm² inside the route where the sounds systems are located. We have good data for crowd densities from our own cordon survey, data on entry/exit volumes at London Underground (subway) stations and 1022 images of the parade taken by police helicopters in the early afternoon of the second day. We show some of these images in Figure 10 from which we have extracted detailed

crowd density-statistics throughout the area. What we do not have are good data on the paths taken by the visitors from their points of entry into the carnival area to the various attractions that comprise the event. However, most visitors enter using one of 38 entry points with half the volume associated with five entry points related to subway stations, and this simplifies the application.

The variant we propose is based largely on Helbing's (1991) social force model, but it also incorporates ideas from fluid flow, queuing theory, event simulation, scheduling and trip accessibility (Still 2001). It is built within a cellular automata structure in which agents embody self-organising behaviour in the form of flocking and swarming (Burstedde et al 2001; Dijkstra et al 2002). In essence, the model generates the relative accessibility of different attractions which make up the entire event with respect to the points where visitors enter and then simulates how these agents walk to the event from these entry locations. In events such as these, we are never in a position to observe the flow of pedestrians in an unobstructed manner because the events are always highly controlled. We have thus designed our model in three stages. First, we build accessibility surfaces from



Figure 10A An ambulance attempting to move injured visitors through the crowd surrounding the parade

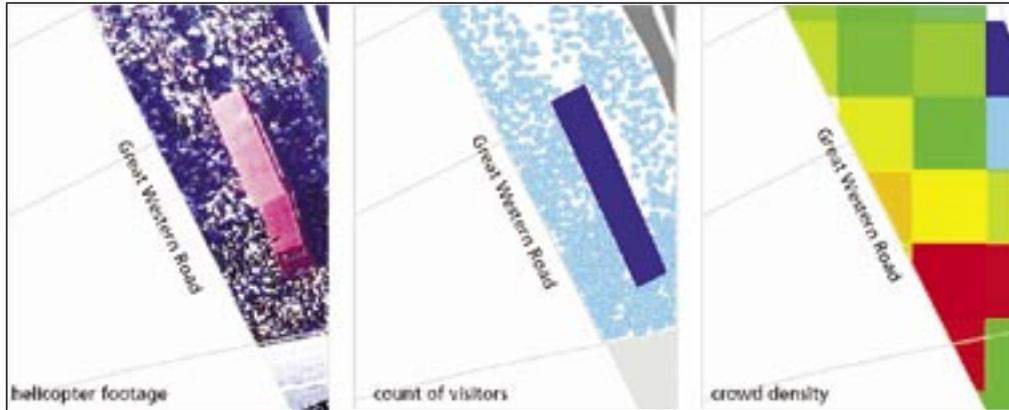


Figure 10B Crowd density at the Notting Hill Carnival - calculating crowd densities from aerial photographs

information inferred about how walkers reach their entry points (origins) relative to their ultimate destinations at the carnival. Second, we use these surfaces to direct how walkers reach the event from their entry points and then assess the crowding that occurs. Finally we introduce controls to reduce crowding, changing the street geometry and volume of walkers entering the event, and fine tune this process iteratively until an acceptable solution is reached. These three stages loosely correspond to exploration, simulation and optimisation.

In the exploratory stage, we begin with walkers located at their ultimate destinations in the carnival area. Walkers move randomly from location to location (which we refer to as cells), avoiding inaccessible cells that are obstacles such as buildings and barriers. The probability of moving from one cell to another is computed according to the accessibility of that cell, which in turn is based on the number of walkers having already visited the cell. At the beginning of this process, all cells have the same probability but eventually a walker will discover an entry point to the carnival—an origin—and when it does, the walker switches from exploratory to discovery mode and returns to the destination with knowledge of the discovery. As the destination is known in that the walker has come from this, it lays a trail back to its source akin to the way ants drop pheromone once they have discovered a food source and head back to their nest (Camazine et al 2001). When the walker enters the neighbourhood of its destination, it switches back to exploration mode and the search begins over again.

It takes some time before agents discover an origin. Before this, the search is a random walk with the route accessibility surface set as a uniform distribution. If a walker crosses the edge of the event space, the walker is absorbed, regenerates at

its destination and begins its search again. In its early stages, this is a random walk with absorbing barriers, much like that shown earlier in Figures 2C and 2D, with the variance of its lengths roughly proportional to the time taken so far to traverse the terrain. As the process continues, more and more origins are discovered, while during exploration, walkers direct their search at routes to origins already discovered. Those origins closest to destinations are discovered first and a hierarchy of 'shortest routes' is built up, continually reinforced by this positive feedback. This is a variant of a generic algorithm predicting trail formation and collective foraging behaviour amongst animal populations. It is extremely efficient for predicting shortest routes in geometrically constrained systems (Bonabeau et al 1999). In Figure 11 we illustrate this for the accessibility surface to various destinations and shortest routes to the tube (subway) stations in Notting Hill. To impress its efficiency, we first show the simulation without obstacles to movement—with the street pattern—in Figures 11C and 11D which shows the true nature of the swarming that takes place, and then with the real street pattern imposed in Figures 11E and 11F.

The exploratory stage finishes when the walkers converge on an efficient and unchanging set of paths between the entry points (origins) and the attractions (destinations) which define the carnival events. In the second stage, we launch walkers from their entry points, and these walkers move towards these events using the accessibility surface based on the shortest paths from stage one to condition the probabilities of movement. There are two effects that complicate this and the first is flocking (social attraction), which directs movement as an average of all the movement in the immediate neighbourhood of a walker (Reynolds 1987, Vicsek et al 1995). But the second effect indicates that a move only takes place if the density of walkers in any cell is less than some threshold based on the accepted standard of two persons/m². If this is exceeded, the walker evaluates the next best direction and, if no movement is possible, remains stationary until the algorithm frees up space on subsequent iterations. These rules are ordered to ensure reasonable walking behaviour. This second stage is terminated when the change in the density of walkers in each cell converges to within some threshold where it is assumed a steady state has emerged.

We can assess how good the model is at predicting the observed distribution of crowds using a battery of statistical tests based on densities and paths and if the model survives, we move to the third stage that is more informal. Note that so far we have not introduced controls on where people are able to move for we begin with no street closures or barriers used in crowd control. In a sense, the carnival



Figure 11A



Figure 11B

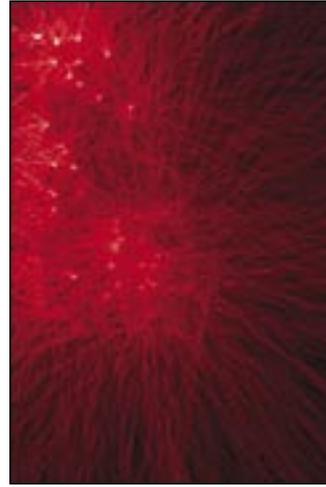


Figure 11C

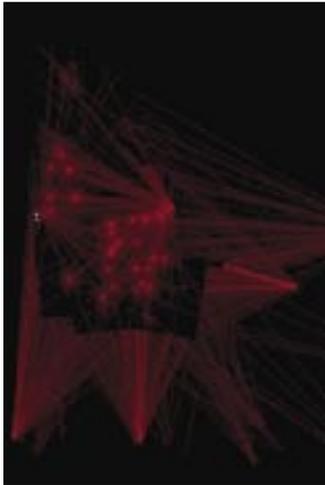


Figure 11D

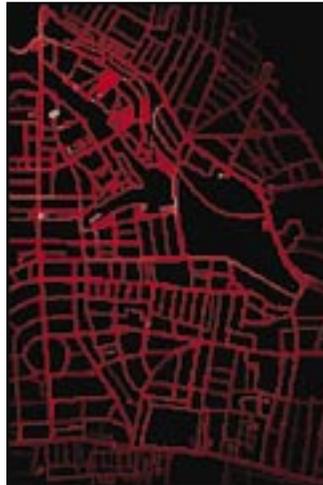


Figure 11E

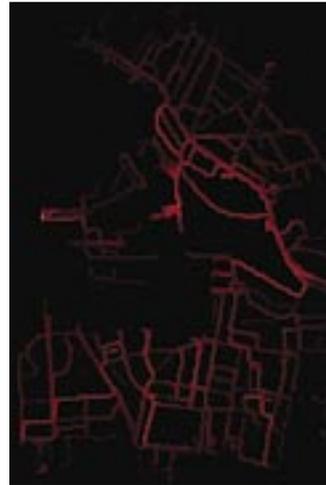


Figure 11F

Figure 11 Exploration of the street system and discovery of entry points (tube stations) in Notting Hill: A, the street geometry; B, the parade route (red), sound systems (yellow) and subway stations (white); C, accessibility from parade and sound systems without streets; D, shortest routes to subway stations without streets; E, accessibility with streets; and F, shortest routes with streets. Relative intensities (of accessibility) are shown on a red scale (light = high; dark = low)

is never without any control so our third stage is to introduce such controls, one by one, to ensure that we produce a safe simulation as well as showing what might be achieved under certain strategies. We do this by examining statistics from the second stage and gradually making changes to reduce the population at risk by introducing barriers, capacitating entry points and closing streets. As the repercussions of this are not immediately obvious, we make these changes one by one, re-running the model until an acceptable solution emerges. In estimation, this stage may also be used to assess the efficacy of existing controls. It is not possible to develop a formal optimisation procedure as so many additional factors such as resources for policing, etc., cannot be embodied in the model. Nevertheless, we consider this interactive method of introducing control the best approach so far for assessing alternative routes.

In implementing the actual model, we start by finding the shortest routes from the parade and static sound systems to the 38 entry points located on the edge of the traffic exclusion zone shown in Figures 12A and 12B. The swarm algorithm predicts the numbers of walkers that ‘find’ each entry point and we compare this uncontrolled prediction to the cordon survey, explaining 64 per cent of the variance. We then use the observed volumes of visitors at entry points to launch these as agents who ‘climb’ the accessibility surface produced in this first stage. We show the crowd density distribution in its steady state in Figure 12C and this identifies significant points of crowding. We predict around 72 per cent of the variance of the observed densities for 120 locations where good data are available. At the third stage, we rerun the model with the official street closures and barriers imposed as illustrated in Figure 12D. At this stage, we have increased the variance explained to 78 per cent, but not all the points of extreme crowding have been removed. This suggests that even in estimating the model, it can be used in a diagnostic manner to identify vulnerable locations as we show in Figures 12E and 12F.

We can now use the model to test alternative routes as we did in the full project (ISP 2002). After considerable political debate, an interim change in route was agreed on for Carnival 2002 where the northernmost section of the parade (shown in Figure 12A in green) was removed. Running the model leads to slightly reduced average crowding, but other problems associated with starting and finishing the parade that are not included in this model emerge. The process of changing the carnival route is still under review with better data being collected each year. Many of the problems of using this model interactively with those who manage

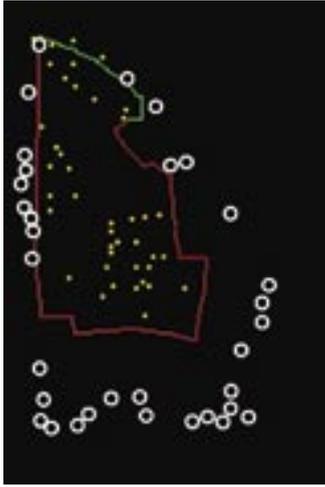


Figure 12A



Figure 12B

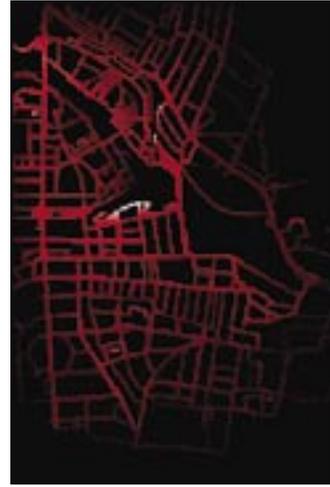


Figure 12C



Figure 12D

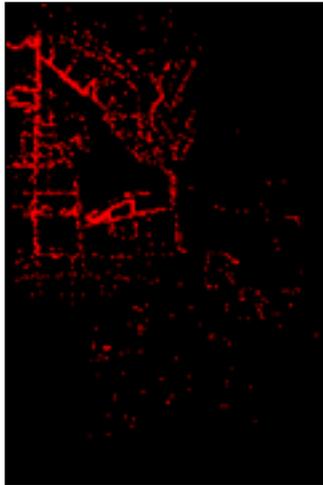


Figure 12E

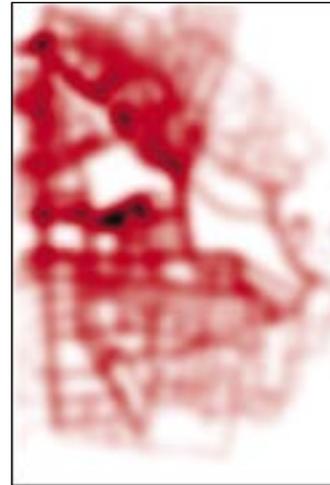


Figure 12F

Figure 12 The full modelling sequence and identification of vulnerable locations: A, the 2001 parade route (red and green) with proposed 2002 route in red, sound systems (yellow), and entry points (white); B, composite accessibility surface from stage one; C, traffic density from stage two; D, areas closed by the police used in stage three; E, location of walkers in the stage three steady state; and F, vulnerability of locations predicted from stage three, on a red scale (lightest red low to darkest red high; the white background are areas not relevant to such vulnerability).

the event are being improved. As we gain more experience with this approach, we are better able to adapt this kind of model to different policy-making situations.

6 The future: fine-scale dynamics and GIS

In this chapter we have shown how a focus on finer scales of granularity than GIS is typically concerned with creates problems of representation and dynamics which are hard to embody within traditional GIS. Clearly, the models we have been dealing with are part of an enhanced geographic information science but a focus on the fine, localised scales with human action and motion at the fore, reveals that traditional GIS is largely concerned with systems represented in terms of their inanimate characteristics. Insofar as human activities are represented in GIS, these are in static, aggregative terms. Moreover, the kinds of events that this chapter has been concerned with are much shorter lived than those that are traditionally a part of our science. This is partly because we are only just beginning to get a handle on fine-scale events of short duration but it is also because such events now seem more important than they ever were, as is witnessed in our concern for developing an appropriate science to deal with safety, crime and leisure.

This chapter also reveals a wider concern with GIScience. Spatial analysis and modelling form very different traditions and, although it is not surprising that much of GIS and GIScience has been developed by those who are involved with these other traditions, it has always been difficult to generate an appropriate fusion of spatial analysis and spatial representation. This has been managed in ad hoc and satisfactory ways for static aggregative phenomena, but once the concern shifts to dynamics and to fine spatial scales, then fusing traditional GIS with these kinds of models and analysis becomes an even greater challenge. The kind of fusion that has been attempted quite successfully by Evans and Steadman (this volume) for land use-transport models is not possible with the models of this chapter because the processes and scales involved are too different from the way traditional software in information systems has evolved. Nevertheless, what we have shown is that the fine scale of granularity and its dynamics are rich in detail and that the visualisation involved in making sense of these events is consistent with the traditions of GIS. We have raised many new issues in this chapter which we see as some of those to which GIS needs to respond to as it becomes increasingly adopted and adapted in urban policy making where physical design and social action are integral to relevant human action.