EG UK Theory and Practice of Computer Graphics (2005)
L. Lever, M. McDerby (Editors)

A Psychologically-Based Simulation of Human Behaviour

Stephen J. Rymill and Neil A. Dodgson

University of Cambridge Computer Laboratory
15 JJ Thomson Avenue, Cambridge CB3 0FD, UK
stephen.rymill@cl.cam.ac.uk, neil.dodgson@cl.cam.ac.uk

Abstract

We describe a system designed to simulate human behaviour in crowds in real-time, concentrating particularly on collision avoidance. The algorithms used are based heavily on psychology research, and the ways this has been used are explained in detail. We argue that this approach gives better results than conventional methods, and detail further work to be done.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Animation and Virtual Reality

1. Introduction

Realistic simulation of a crowd of people is a challenging area of computer graphics. Many of the problems with creating lifelike 3D animated models have been solved, but the difficulty now lies in creating behaviour that is believable. In this paper we show that the quality of the behaviour can be improved by making use of theories from psychology. The goal of our work is to produce animations of a virtual environment populated with virtual actors interacting with their surroundings and with each other.

This paper presents a system that simulates the behaviour of a large virtual crowd in real-time. Evidence from the psychology literature, together with the results of observations, is given to demonstrate the foundation for the behaviour techniques chosen. Methods of collision avoidance are described, in particular a novel way of avoiding an oncoming actor and new techniques for overtaking. There is also a brief discussion of ways that behaviour can be varied, based on characteristics assigned to an actor.

There are a number of applications for this type of simulation. One of these is in architecture: for example, the design of a new building can be tested to see how efficiently it can be evacuated, or the effects of more turnstiles at a stadium can be tested. For these uses it is clear that an accurate simulation of behaviour is needed, otherwise the testing will be of little use. Other uses are for the creation of visual effects for movies, where expensive crowds of extras can be replaced by a virtual crowd. A real-time simulation of a crowd could also be used in computer games, in which the actions of the player may have an effect on the crowd's behaviour. In both of these applications, the requirement of absolute accuracy is not as strict, but the behaviour must still look believable to the viewer.

Section 2 covers previous work in the field of behaviour simulation. Section 3 gives details of the psychology literature used. Section 4 describes the algorithms that have been implemented. These are then evaluated in section 5, and discussed in section 6. Conclusions, along with ideas for future work, are given in section 7.

2. Previous Work

The simulation of behaviour has been studied since the earliest days of computer graphics research. Early work concentrated on animal behaviour, with birds a popular choice, but recently there has been a lot of work on human behaviour. Techniques for simulating a crowd as a single entity have been proposed, as well as those which consider each person in the crowd separately.

The work which popularized the whole field was carried out by Reynolds [Rey87] in 1987. He produced a simulation of a flock of birds, which he named boids. The key point that he introduced was that each boid should be treated separately; the combination of each boid following a standard set of rules would result in a flocking behaviour for the entire group. Each boid was given three principal rules: move towards other boids, try to match velocity with nearby boids, and avoid collisions. This work was continued by Tu and Terzopoulos [TT94], in a very detailed simulation of a school of fish, using the same principles as Reynolds. They went further by including a complicated system to deal with goals, giving each fish a mental state and working out their intentions from this. Reynolds continued his work [Rey99], building on his original model to give an implementation of some simple steering behaviours, such as seek, pursuit and obstacle avoidance; combined together these can produce fairly complex behaviour.

Turning to the simulation of people, a number of researchers have attempted to simulate the behaviour of a crowd as an entity in its own right. Thalmann, Musse and Kallmann [TMK00] present a system which collects actors together if they share a common goal and controls them as one group. This is based on work by Musse and Thalmann [MT97], which included a multiresolution collision avoidance scheme. In his PhD thesis, Still [Sti00] presents a design for a system based on watching real crowds at football stadiums; this is a rare example of a system based on observations. His approach simulates the crowd as an emergent phenomenon using simulated annealing and mobile cellular automata. However, the problem with this emphasis on simulating the crowd as a whole is that it goes against Reynolds' idea that the behaviour of the group should be built up from the actions of all its members. While the overall behaviour of the crowd may be realistic, closer inspection may show that the individuals are not behaving correctly.

There has also been work in the specific area of simulation of individual human behaviour. Thompson and Marchant [TM95] created a simulation for the fire safety community to assist in designing buildings so they can be evacuated quickly. Before the simulation is started the optimal escape route is calculated for every position within the building; the actors have access to this optimal route during the simulation, and so always know the best way to go from their current position. Actors can overtake each other, and will do so if it will increase their speed; they will choose which side to overtake based on which deviates less from the way out. The speed of an actor is determined from the distance to the nearest person using a simple relationship. All of these rules are rather unrealistic, and are not based on the psychology of how people actually behave. In particular the knowledge of the optimal escape route has a very detrimental effect on the realism of the simulation. However, the system proved sufficient for use in the very restricted scenario for which it was designed.

Other work in this area is by Feurtey [Feu00], who uses a spacetime approach to predict collisions with other actors, Helbing and Molnar [HM95], who use a social force model to simulate movement based on motivations, Blue and Adler [BA00], who use a cellular automata model, Gillies and Dodgson [GD], who concentrated on obstacle avoidance and the simulation of attention, and Lamarche and Donikian [LD04], whose work on path finding also includes ideas on behaviour simulation. Recently the commercial product Massive [Mas03], made famous by *The Lord of the Rings* film trilogy, has been released. It has been designed to be able to simulate crowds in the hundreds of thousands, and uses a neural network combined with fuzzy logic. The brain of each actor consists of up to 2000 nodes, and takes as input the data from sensors recreating eyes and ears; the output is a command to the rendering engine for which piece of motion capture data to use. While producing impressive results, the creator of Massive is quick to point out that it is not an artificial intelligence system. Anecdotal evidence is that Massive simulations need to be carefully set up, and that the simulations can run for only about thirty seconds before unrealistic or unwanted behaviour starts to become obvious.

3. Psychology Background

The aim of this work is to produce a simulation based on real human behaviour, so the psychology literature was studied to find theories that attempt to explain behaviour, and that would be suitable for conversion to algorithms. The main sources are works by Wolff [Wol73], Collett and Marsh [CM81] and Goffman [Gof72]. Relevant results from these are presented below, divided into categories based on the aspect of behaviour in question. Most of the material relates to walking on a pavement, where there is restricted space, but it can also be applied to an open area.

3.1. Avoiding oncoming people

1. In experiments carried out by Wolff, he observed that, in high density crowd situations, pedestrians would start to avoid an oncoming collision late on, when the other was 5 feet (1.5m) away or nearer. In low density cases, with only a few people around, this would rise to much greater distances – up to 100 feet (30m). Goffman agrees that people will tend to avoid each other as early as possible, so that only a small angular direction change is required; this is thought to be explained by a desire to maintain dignity. Higher density crowds clearly offer less opportunity for early avoidance.
2. Wolff found that the avoidance behaviour used by pedestrians varied based on the density of the crowd. In low density areas, where a collision can be predicted well in advance, pedestrians would change their paths and detour around the oncoming person (who may also do the same). However, at higher densities, when a collision is not predicted until late on, a common behaviour, occurring about half the time "was not total detour and avoidance of contact, but a slight angling of the body, a turning of the shoulder and an almost imperceptible side step." Wolff calls this a 'step-and-slide'. He also notes that both people in the encounter have to perform this manoeuvre

as “if either of the two had walked directly ahead, there would have been a collision at the next step."

3. Observations by Collett and Marsh indicated that there are two types of step-and-slide: an ‘open’ pass, where the pedestrian turns his body towards the oncomer, and a ‘closed’ pass, in which he turns away. They also found that women prefer the closed pass and men the open.
4. Collett and Marsh also observed the direction of the detour made around an oncoming pedestrian. In the cases where the pedestrians’ paths did not completely overlap, the direction of detour was always determined by the relative positions of the two participants; for example, if their right shoulders were overlapping, they would both move to the left. However, if their paths overlapped totally then a choice clearly had to be made; results showed that twice as many pairs moved to the right as to the left. This fits in with Goffman’s statement that pedestrians tend to avoid collisions by veering to the right, possibly explained by the fact that the majority of people are right handed.

3.2. Avoiding people moving in the same direction

If a collision is predicted with a pedestrian ahead of the subject there are two options: overtake, or slow down and walk behind them.

1. Wolff observed that people do not position themselves directly behind the pedestrian in front of them, but behind their shoulder; this is so that they can still monitor conditions ahead. The only exception to this rule is when the gap between people is more than 15 feet (5m), in which case the following actor will walk straight behind the leader.
2. Wolff further observed that pedestrians form clusters, based on their speed of travel; these clusters grow by aggregation from the rear. At high crowd densities, more clusters form, and they grow bigger; this is because, in a more crowded situation, there is less space to overtake, so slowing down is an easier option.
3. Collett and Marsh also studied the idea of clusters, which they think of as lanes of people. They suggest that in high density crowds there is more advantage to lane formation, with groups of people all walking at the same speed, as less effort is then required to avoid collisions.
4. If the decision is made to overtake then, according to Collett and Marsh, pedestrians can weave their way through their own lane of people, or make use of the interface between opposing lines of people, or step off the pavement onto the road. Wolff also notes that if a pedestrian in a cluster wishes to travel faster, then they can either thread their way through the cluster, or go around it. Goffman mentions that one can almost always overtake if desired, even if it requires going onto the road to do so; however, there are cases where it is impossible, in very dense crowds, or in restricted spaces.

3.3. Resuming original course

1. Both Goffman and Wolff agree that, after detouring to avoid a collision, rather than going straight towards their goal, people will return to the original path (or line-of-walk) that they had been on before the detour. The explanation for this given by Goffman is that in their original path they have already conducted long-range scanning of oncoming pedestrians, which would need to be done again if they proceeded on a new path. An alternative theory suggested by Wolff is simply that people have a preference for which part of the pavement they like to walk on – for example, some people may prefer to walk on the innermost part of the pavement; they will therefore return to this preferred path after the detour.
2. Wolff notes one exception to the above rule: if the subject were about to turn off the pavement anyway, for example to enter a shop, then, after avoiding the collision, they would walk straight towards their goal, rather than returning to their original path.

4. Behavioural Algorithms

The algorithms used in our work to control the behaviour of the actors can be divided into three main categories: predicting when collisions will happen, avoiding those collisions, and finally returning to their original path.

4.1. Collision prediction

Each frame of the simulation, every actor needs to check for future collisions with all other actors in the scene. Because this would otherwise be a slow process, the database of actors is filtered, so that only those actors within the field of view of the subject, and within 30 metres, are investigated. This idea was used by Gillies [GD] and reflects the fact that people only pay attention to that which is important to their current situation. To be in the field of view of the subject, an actor must be within 60° either side of the direction in which the subject is travelling. 30 metres corresponds to the maximum distance at which collision avoidance would occur (section 3.1.1.); beyond this range there is no need to predict that the collision will happen.

Each actor is represented by a centre, a radius and a velocity. The circle defined by these represents the personal space around the actor; it is centred on the actor’s face, as the space should extend further in front of the actor than behind.

For each pair of subject and other actor, their relative position after t frames is as shown in the following equation:

$$\mathbf{p}(t) = \mathbf{p}(0) + \mathbf{v}t \tag{1}$$

where $\mathbf{p}(t)$ is their relative position at time t, and $\mathbf{v}$ is their relative velocity. From this the distance between them at time t, $\mathbf{d}(t)$, can be found using:

$$|\mathbf{d}(t)|^2 = |\mathbf{d}(0)|^2 + 2t(\mathbf{v}\cdot\mathbf{d}(0)) + t^2|\mathbf{v}|^2 \qquad (2)$$

Therefore to find if a collision will occur one must check if this distance is ever less than twice the radius of an actor (r); this is done by solving the following equation for t:

$$|\mathbf{d}(0)|^2 + 2t(\mathbf{v}\cdot\mathbf{d}(0)) + t^2|\mathbf{v}|^2 = (2r)^2 \qquad (3)$$

There can be up to two solutions, t_1 and t_2, with $t_1 < t_2$. The only interesting solutions are for $t_1 \geq 0$, which is a collision that will happen after t_1 frames, and $t_1 < 0$, $t_2 > 0$, which is an ongoing collision.

Each actor stores a list of predicted collisions, together with the time remaining until each collision would occur. These are sorted so that the nearest collision can be given the highest priority.

If a collision has been predicted from equation 3, then the type of collision must be determined. There are three possible types of collision, which here will be called Towards, Away and Glancing; these are shown in figure 1. If $\mathbf{v}_s$ is the velocity of the subject, and $\mathbf{v}_c$ is the velocity of the actor that will collide with it – its 'collidee' – then:

- A Towards collision occurs if the actors are walking towards each other, and will happen if $\mathbf{v}_s \cdot \mathbf{v}_c < 0$.
- An Away collision is, conceptually, when the subject is behind the collidee, but is gaining on them. To check for this, we find a vector perpendicular to the vector between the two actors. If the directions of $\mathbf{v}_s$ and $\mathbf{v}_c$ are on the same side of this vector, it is an Away collision.
- A Glancing collision is a side-on collision between two actors walking in roughly the same direction, and will occur if $\mathbf{v}_s$ and $\mathbf{v}_c$ are on opposite sides of the perpendicular vector described above.

If a collision is predicted then a further check is made to see if the collidee is part of a group of actors. Any actor within a small distance of the collidee, and travelling in roughly the same direction, such that there would not be room for the subject to walk between them, is defined as being in a group with the collidee. Details of which actors are in which groups are stored, and group membership is checked every frame, to ensure actors are still in the same groups. Actors do not need to know that they are in groups, as all the information is calculated automatically. When considering avoidance strategies the actor will look at the group as a whole, if possible avoiding the whole group rather than just the collidee.

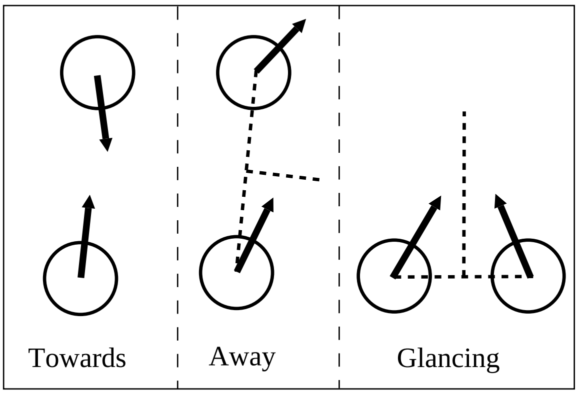

Figure 1: *Three different types of collision*

4.2. Collision avoidance

If a collision has been predicted then an attempt must be made to manoeuvre away from the collision. This attempt starts as soon as the prediction has been made, in accordance with the theory that people try to maintain dignity by making the smallest detour possible (section 3.1.1.). The behaviour is different depending on the type of the collision, and on the time until the collision is due to occur. Actors principally try to avoid the nearest collision, but in doing so they will try not to cause collisions with other actors. If only the nearest collision was considered then a situation could arise where an actor oscillates between different courses as different collisions become the nearest; this would look unrealistic.

4.2.1. Towards collisions

The first stage is to determine whether the collidee is to the left or right of the subject, by considering their relative positions and velocities. As mentioned in section 3.1.4., people will prefer to pass on the side with least deviation from their path: if their right shoulders are overlapping – meaning that the collidee is to the right of the subject – the subject should detour by moving to the left. Priority is therefore given to this behaviour, but other factors may mean the actor will detour on the other side.

Observations show that the subject has three different ways of avoiding the collision: changing direction only, changing speed only, or changing both direction and speed. Each technique is tried, until one is found that works; therefore they are ordered in terms of increasing disruption to the subject's walking. Changing direction is less disruptive than changing speed, according to Goffman. This is explained by pedestrians having already chosen the speed they can comfortably walk at: walking faster will require more effort, while walking slower will hinder their progress.

1. Changing direction. Direction changes of increasing

magnitude are tested, starting with a detour around the shortest side, and then trying the other side. If there is total overlap between the two actors – they are walking straight towards each other – then the right side is tried first, in accordance with observations (section 3.1.4.). Six different direction changes are tried on each side, to give an easy-to-calculate sample of the continuous range of options a real person would consider. Each of the 12 direction changes are tested to see if walking in that direction would cause a collision to happen, either with the current collidee, or with any other nearby actor. If a direction change is found that would avoid these collisions then it is chosen, otherwise the next technique is tried.
2. Changing speed. Speed changes of up to 50% slower and 50% faster are tested, starting with small changes and working up to larger changes, in increments of 10% (another discrete representation of the real-life continuous choices). If any of these speed changes produces a path with no collisions with nearby actors, it is selected as the avoidance behaviour.
3. Changing direction and speed. If no suitable behaviour has yet been found, all combinations of direction change and speed change, used above in isolation, are tested. This is left until last, as such behaviour will cause significant disruption to the subject's route. As before, detouring on the shortest side is preferred.

If an avoidance behaviour involving a direction change has been selected, an appropriate vector is added to the subject's velocity; this is designed to simulate the way that people do not actually turn onto a new course to avoid a collision, but just introduce some sideways movement into their steps.

One further technique is available to actors: if they have managed to avoid a direct collision, but their paths still overlap, they can perform the step-and-slide manoeuvre described by Wolff (section 3.1.2.). Both actors will have predicted the collision, so both will perform the manoeuvre, making it more effective. It is carried out by the actors angling their bodies away from each other, while continuing to walk forwards; this reduces the overlap of their bodies, so that they can pass without touching. It takes a few frames for the shoulder to swing around in the course of the manoeuvre, so the action needs to be started just before the collision would occur. It should be noted that during the manoeuvre the actors will pass through each other's personal space; this is acceptable behaviour as long as the bodies do not collide.

The step-and-slide technique is very useful in collisions which have been predicted late on, when simple detouring would not take effect quickly enough; it is therefore an example of how different techniques are used depending on the estimated time to collision.

If no behaviour has been found that will avoid a collision, as a last resort the subject simply stops walking. This will then allow the other actor involved in the collision to avoid the subject, who can then resume walking. One problem with this technique is that it can occasionally lead to a deadlock where none of the actors can proceed. However, the situations where this can happen are rare, and exist in real life too, where complicated behavioural rules are needed to resolve them. Such rules are beyond the scope of the present work.

4.2.2. Away collisions

An Away collision is one where the collidee is in front of the subject, but the subject is walking faster than the collidee, and so would bump into the rear of the collidee. To deal with this situation, the subject has two choices: slow down to the same speed as the collidee and walk behind it, or overtake. For the purposes of these experiments, the decision on which to do is made using a simple mental model: each actor is given a variable indicating the degree to which it is willing to slow down, which is a representation of whether the actor is in a hurry.

If the decision is made to slow down, the speed of the subject is set to be the same as the collidee. This speed will then be preserved until there is a change in conditions, for example if the collidee moves out of the way.

The procedure for overtaking is more involved. First it is determined which side of the subject the collidee is on; this is so that the overtaking can be done on the shortest side, if possible. A check of the entire potential route of the subject during overtaking is then made, to see if the procedure could be conducted without colliding with other actors. If this is not possible, then the actor temporarily slows down behind the collidee, and waits until there is space to overtake. Once there is space, the actor's speed is boosted by 25%: this is an estimate based on the observation that people generally walk faster when overtaking, to allow the procedure to be finished quicker. A sideways component is added into the subject's velocity, to actually manoeuvre around the collidee.

4.2.3. Glancing collisions

Glancing collisions are dealt with in a similar way to Towards collisions. The subject tries a variety of direction and speed changes, in an attempt to find a course that will avoid a collision occurring. If no such course can be found, the subject is forced to stop until the collidee has passed.

4.3. Returning to path

As discussed in section 3.3, both Goffman and Wolff state that after an actor has successfully avoided a collision, it should return to its original path. The first problem is to determine when collision avoidance is actually over:

- For a Towards or Glancing collision, a collision can be over in two ways: either the subject is now in front of the collidee, or it is a significant distance away.

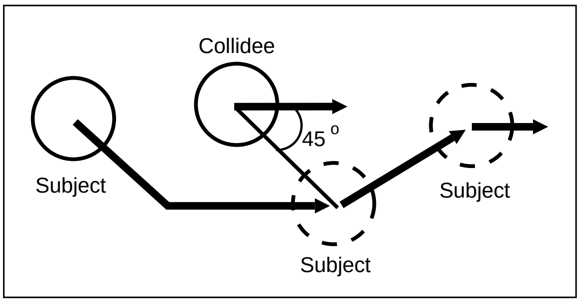

Figure 2: *Overtaking theory*

- For an Away collision, when the subject is overtaking the collidee, the subject needs to go a small distance past the collidee before cutting back in front of it, otherwise a glancing collision will occur. Therefore the angle between the collidee's velocity and the line joining the actor's centres is found; if it is less than 45°, and the subject is in front of the collidee, it is safe for the subject to return to its original path. This is shown in figure 2.

Each actor has a goal that it is is trying to reach by following a path represented by an ordered list of nodes. The actor should return to this path, but it should also not suddenly change direction to do so, as this would look unnatural. So the path ahead is searched to look for the first node which would require a direction shift of less than 45°. If one is found then a course is set for this node; once it has reached this node, normal following of the path is resumed. If no such node is found then a new route is calculated towards the actor's goal.

The exception mentioned in section 3.3.2. has also been implemented: if the actor is very close to its goal after the collision then it will walk straight towards it, rather than first going back to its original path.

5. Results

The behavioural algorithms described above have been implemented in a system called CrowdSim, using C♯ and DirectX to provide a simple 3D representation of the actors. The user interface has been designed to allow easy testing: the simulation can be paused at any time, and replayed at any speed, or frame by frame, meaning that a collision can be viewed from a variety of angles; information can also be displayed about any of the actors just by clicking on them.

Although the system can simulate any type of open or enclosed space full of people, the examples shown here are of a typical street scene, with pavements and a pedestrian crossing. The narrow space of the pavements produces a large number of potential collisions, and the pedestrian crossing acts as a bottle-neck; together these produce a challenging, and interesting, situation. It is also a familiar situation, so it is easy for a viewer to tell if it looks realistic.

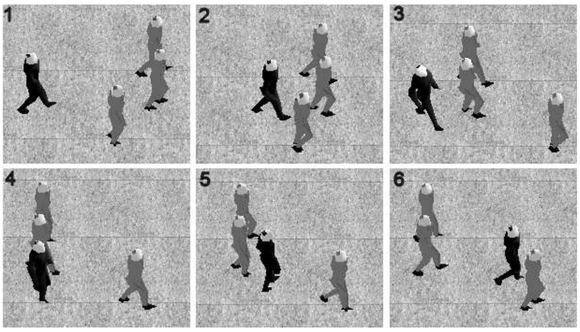

Figure 3: *Six frames from points in the overtaking procedure: in frame 1, the subject approaches three other actors; in frame 2 it slows down as there is no space to overtake; in frame 3 there is now space to overtake so it begins to do so; frame 4 is during the overtaking process; in frame 5 the danger of collision is now over so it can return to its original path; by frame 6 it is back on its original path.*

This system is designed to be used in real-time, so that different ideas can be tested quickly. On a fairly high-end PC (AMD 2700+) a crowd of 500 people can be simulated at an acceptable frame rate (20 frames per second) in the street scene shown. In scenes where the actors are more spread out, larger crowds – up to around 1000 – can be simulated in real-time as there will be fewer potential actor interactions.

Figures 3-5 demonstrate the system in action, highlighting the overtaking procedure and lane formation.

6. Discussion

The system incorporates a number of features that make it more realistic than existing simulations:

- The step-and-slide manoeuvre, a pervasive feature of everyday walking, which has never before been simulated.
- The procedure involved in deciding whether to overtake, and when it is possible to do so. Overtaking was considered by Thompson and Marchant [TM95] (section 2), but in their system actors would always overtake if they were going faster than their collidee, and the side was chosen simply as the one deviating least from the way out. Lamarche and Donikian [LD04] consider simple overtaking, but not the issue of deciding when to overtake.
- The return-to-path behaviour has not been simulated before: in previous work, after the collision has been avoided the actor either just continues walking in that direction, or plots a new course to its goal from that position.

It is our belief that aspects of our simulation compare favourably to the real world. Personal observations of busy pavements in shopping streets have shown that the step-and-slide is used, in a variety of forms, extremely frequently. The overtaking procedure implemented here is also a very good

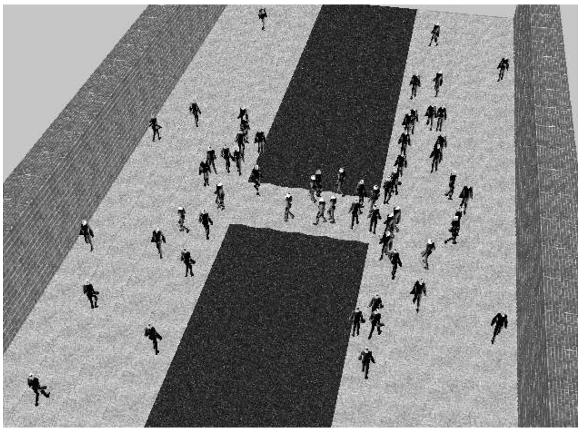

Figure 4: *A typical street scene, with actors spread out over two sections of pavement separated by a road*

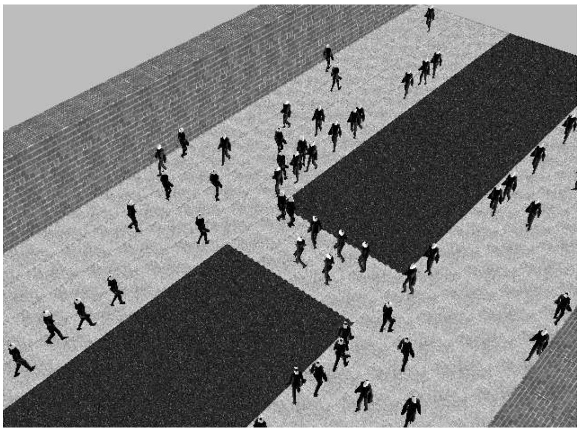

Figure 5: *Examples of lane formation. Lanes can be seen along the edges of the road, and on the crossing*

representation of observed behaviour. The actors in the simulation can be seen to form lanes, moving in the same direction at the same speed; this matches the observations of Collett and Marsh, and Wolff (section 3.2). Oncoming collisions are also handled well, with the actors passing on the predicted side.

7. Conclusions and Future Work

This paper has presented useful steps towards a system that can accurately simulate human behaviour based on theories from psychology. The ideas discussed here mainly centre on techniques for collision avoidance. These techniques add to the realism of the simulation. This realism is achieved efficiently: the system can simulate a large number of actors in real-time.

There are many ideas for future work to improve the realism of the simulation, in the following areas:

1. Actors are currently able to see everybody within their field of view, and access accurate information about their positions and velocities. This does not represent the real world, where each pedestrian has to build up for themselves a mental model of where other people are, and how fast they are travelling; this model is assembled and then maintained as the pedestrian's gaze moves around the scene. Because of this, the model will not be completely accurate or up to date; when simulated this should have an observable effect on the actor's behaviour. One success criterion could be having two actors collide if they are not paying sufficient attention to the world around them, something which does happen in reality. Gillies [GD] undertook some excellent work on attention and we are incorporating some of his ideas, in addition to our own, into our work.
2. Currently each actor is assigned a random goal to walk towards; when it gets there another random goal is chosen. A better system would involve the actor being given a plan, which would consist of multiple goals. For example, an actor could be told to stroll around a park, or visit a shopping mall, and suitable goals would be created automatically. We thus would have one mechanism for performing high-level route planning and another, described in this paper, for making real-time adjustments to the route to avoid collisions.
3. There is scope for more work on the actor's mental model, currently used to rate the actor's urgency in getting to its goal and therefore how much they are prepared to slow down rather than overtake. This could also affect how willing an actor is to detour to avoid a collision, how much attention they pay to others, whether they are willing to split from a group of people, and how willing they are to take risks such as walking on the road.
4. Currently all actors walk through the scene on their own; further work would allow actors to walk in pairs or small groups.

References

[BA00] Blue V., Adler J.: Cellular automata model of emergent collective bi-directional pedestrian dynamics. In *Artificial Life VII, The Seventh International Conference on the Simulation and Synthesis of Living Systems* (2000). 2

[CM81] Collett P., Marsh P.: Patterns of public behavior: Collision avoidance on a pedestrian crossing. *Non-verbal Communication; interaction and gesture, Approaches to Semiotics* (1981), 199–217. 2

[Feu00] Feurtey F.: *Simulating the Collision Avoidance Behavior of Pedestrians.* Master's thesis, Department of Electronic Engineering, University of Tokyo, 2000. 2

[GD] GILLIES M. F. P., DODGSON N. A.: Attention based obstacle avoidance for animated characters. In *Virtual Reality*. To appear. 2, 3, 7

[Gof72] GOFFMAN E.: *Relations in Public*. Pelican, 1972. 2

[HM95] HELBING D., MOLNAR P.: Social force model for pedestrian dynamics. *Physical Review E51* (1995), 4282–4286. 2

[LD04] LAMARCHE F., DONIKIAN S.: Crowd of virtual humans: a new approach for real time navigation in complex and structured environments. In *Computer Graphics Forum* (September 2004), vol. 23, pp. 509–518. 2, 6

[Mas03] Massive Limited, www.massivesoftware.com, 2003. 2

[MT97] MUSSE S. R., THALMANN D.: A model of human crowd behavior: Group inter-relationship and collision detection analysis. In *Computer Animation and Simulations, Proc. Eurographics workshop* (1997). 2

[Rey87] REYNOLDS C. W.: Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques* (1987), ACM Press, pp. 25–34. 1

[Rey99] REYNOLDS C. W.: Steering behaviors for autonomous characters. In *Game Developers Conference* (1999). 2

[Sti00] STILL G. K.: *Crowd Dynamics*. PhD thesis, Mathematics Department, Warwick University, 2000. 2

[TM95] THOMPSON P., MARCHANT E.: A computer model for the evacuation of large building populations. In *Fire Safety Journal 24* (1995), pp. 131–148. 2, 6

[TMK00] THALMANN D., MUSSE S. R., KALLMANN M.: From individual human agents to crowds. In *Informatik/Informatique - Revue des organizations suisses d'informatique* (2000). 2

[TT94] TU X., TERZOPOULOS D.: Artificial fishes: physics, locomotion, perception, behavior. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques* (1994), ACM Press, pp. 43–50. 2

[Wol73] WOLFF M.: Notes on the behaviour of pedestrians. In *People in Places: The Sociology of the Familiar*, Birenbaum A., Sagar E., (Eds.). New York: Praeger, 1973, pp. 35–48. 2