Intuitive Method for Pedestrians Simulation

Jérémy Boes, Stéphane Sanchez, Cédric Sanza
IRIT, Vortex team
Université de Toulouse, 31062 Toulouse, France
boes@irit.fr, sanchez@irit.fr, sanza@irit.fr

Abstract

Recent works about pedestrian simulation can actually be sorted in two categories. The first ones focusing on large crowd simulation aim to solve performance and scalability issues at the expense of behavioral realism of each simulated individual. The second ones aim at individual behavioral realism but the computational cost is too expensive to simulate crowds.

In this paper, we propose an alternate approach combining a light reactive behavior with cognitive strategies issued from real life videos. This approach aims at the real time simulation of small crowds of pedestrians (one to two hundred individuals) but with concerns for visual realism regarding heterogeneous behaviors, trajectories and positioning on sidewalks.

1. Introduction

There are many approaches to pedestrian simulation. Discrete crowds (also called agent-based simulations) focus on individuals. Local behavioral rules are given to each agent and a realistic global behavior is expected to emerge [Shao06], [Pelechano07]. Since each agent makes its own decision, discrete crowds allow a great diversity among pedestrians and provide very realistic behaviors, but the computational cost is expensive, which limits the size of the crowd it can handle. Another drawback of discrete crowds is that the local rules are difficult to create for several rules are needed to achieve simple tasks like obstacle avoidance. Since agents have to perceive the world they populate, discrete crowds are dependent of the type of the environment (indoor or outdoor) and of the way it is constructed.

Continuous crowds have a global point of view. Crowd motion is computed with a potential field that every pedestrian follows [Hughes03], [Treuille06]. Recently, aggregate dynamics have combined discrete and continuous models to reach a large number of pedestrians and to handle very dense crowds [Narain09]. Continuous and aggregate crowds can deal with large dense crowds but are less realistic than agent-based simulations, especially when our eye is caught by one particular character in the simulation. These three approaches are the most popular, but not the only ones existing. The crowd patches method puts together patches of precomputed trajectories [Yersin09], it allows to populate infinite worlds but virtual humans are not autonomous and the simulation lacks of interactivity. The crowd by example method constructs a database of situations from the tracking of videos of real crowds. Virtual pedestrians search the database to copy the appropriate trajectory [Lerner07]. This method shows very realistic behaviors, but for a small number of pedestrians.

Ennis et. al have led a study to find out which criteria are determinant to make a realistic simulation [Ennis10]. Most of these criteria are already taken in account by existing simulations, like obstacle avoidance and walking in appropriate areas, but one is not: walking in small groups. Despite it is not mentioned by Ennis, having pedestrians with heterogeneous appearances and behaviors is an obvious key to realism, but not often present in existing models.

In this paper, we present an intuitive approach based on real life observations of pedestrians. The proposed method combines a reactive algorithm of collision avoidance and behavioral strategies. Our goal is to improve visual realism by simulating heterogeneous behaviors and by maintaining small groups of pedestrians.

The second section presents the model of pedestrian. The third section shows the results by comparing the simulation and the real world. The next section concludes this paper and gives further works.

2. Approach

We focused on three main goals. First, simplicity and genericity: we wanted our method to be easily implemented in any environment. We also wanted it to be able to allow heterogeneous behaviors and small groups of people. To help us to reach a great degree of realism we shot videos of pedestrians walking the downtown streets of Toulouse, France. We extracted precious data from these videos. They are described later. Section 2.1 explains how we performed the classic tasks of collision avoidance and retention in walkable areas. Sections 2.2 and 2.3 present how we introduced heterogeneity and small groups.

2.1 Collision avoidance and walkable areas

2.1.1 Obstacle Avoidance

We observed on the videos that pedestrians seem to follow a free space created by those who precede him: they favor *free directions.* We represent these *free directions* in a simple table, called the *direction table*. Each agent has its own table and each cell corresponds to a direction it can take. At any time of the simulation, the value of the cell is the distance the agent can walk, following the corresponding direction, without running into an obstacle (figure 1). The table is initialized with a value that is the maximal distance for which obstacles are taken in account. When an obstacle is perceived, the distance to this obstacle is inserted in the appropriate cell only if it is inferior to the current value of the cell.

Each agent computes its *desired direction*, it is the direction it needs to adopt in order to reach its target. Once its *direction table* is updated according to its perceptions, an agent checks the table if its *desired direction* is free. If not, the agent will have to look for the *closest free direction*, it is the nearest cell containing the maximal distance. The final direction that the pedestrian takes is a weighted average of his

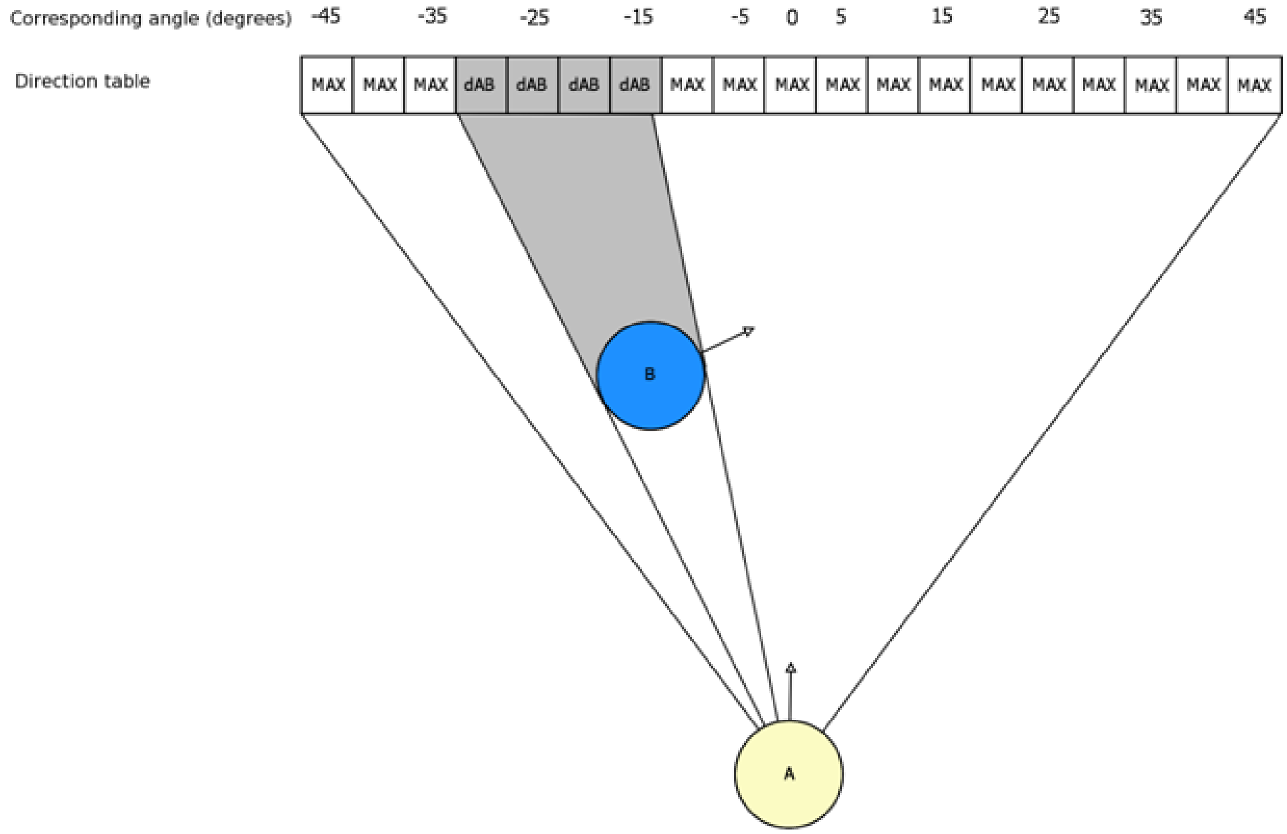

Fig. 1 - Agent A perceives agent B on the left. Some cells of the first half of the table are filled with the distance between A and B, the others are "empty", they contain the maximal distance. Arrows represent the current direction of the agent.

desired direction and the *closest free direction*, with a greater weight for the latter.
If two cells can pretend to be the closest free direction, the cell with the greater index is chosen. This simulates the natural tendency of people to avoid an obstacle by the right rather than the left when the two solutions are equivalent. In order to obtain smoother trajectories, an agent can adjust its direction even if its *desired direction* is free. This happens when an adjacent cell of the one corresponding to the desired direction contains a small distance. This means that an obstacle is near the trajectory, the agent will then shift its orientation from a cell on the other side in order to not get too close of the obstacle.
The number of cells depends on the angle pitch between each cell. If the pitch is too small, agents don't modify their trajectory strongly enough, if it is too high agents shake and have unnatural trajectories. A pitch of five degrees proved to be the best compromise.

2.1.2 Anticipation

Most of the collisions are easily avoided with this technique (especially with static obstacles), but some still occasionally occur with moving objects. To prevent such collisions, agents don't only perceive size and position of other objects, they also perceive speed and orientation. Therefore they are able to extrapolate the trajectory of other agents. The anticipated position (and not the current position) of perceived obstacles is used to update the *direction table* (figure 2). The amount of time over which the agent anticipates depends on several criteria: its speed (more he is fast, less he anticipates over a long time), the distance to the other agent (almost no anticipation for very close obstacles) and the angle between the two trajectories (maximum anticipation for perpendicular trajectories, almost no anticipation for parallel trajectories).
Static and moving obstacles are avoided thanks

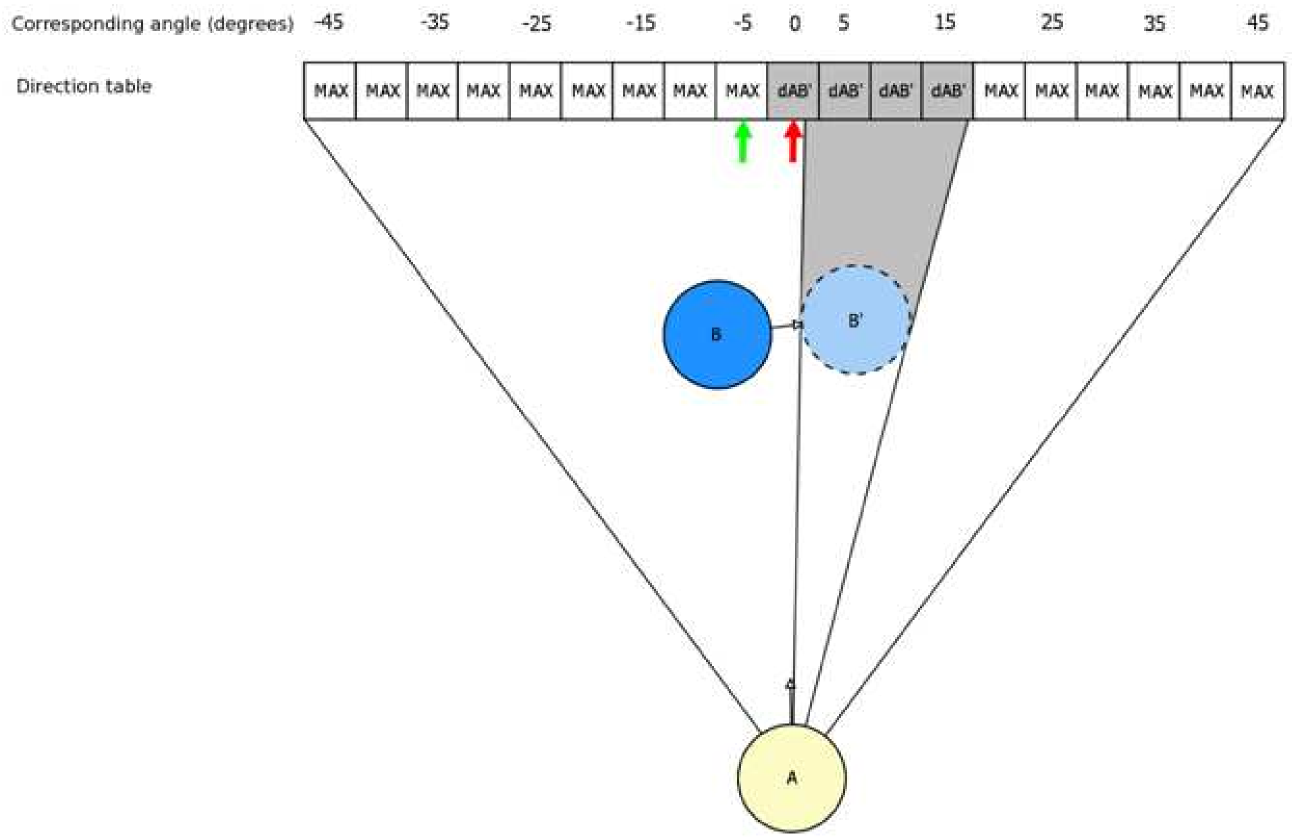

Fig. 2 - Agent A extrapolates agent B position. A will avoid B by the left. Without anticipation, A would have turn right and a collision would have occurred as B is moving this way. Green arrow show the closest free direction cell, red arrow the desired direction cell.

to the same technique, using the same direction table. This allows our method to be easily implemented on any environment: the only requirement is the perception of distance, position, size, speed and orientation, which is basic. Moreover this technique sticks to reality: if an obstacle stands in our way, we adjust our trajectory just enough to avoid it.

2.1.3 Walkable Areas

To ensure that agents stay on the pedestrian network (sidewalks and crosswalks), we tagged borders with *border cells (figure 3)*. They are perceived by agents and treated as obstacles by the direction table. Agents tend to avoid borders, and stay in safe zones. Border cells are not physical obstacles, if an agent is pushed through a border (it happens when sidewalks are crowded), he will cross it and walk on the road. The direction table allows agents to slightly adjust their trajectory but not to make brutal changes, therefore if a pedestrian walks quickly perpendicularly to a border (it happens if his target is on the road), he will cross it. Of course the treatment of border cells is deactivated for pedestrians who walk on the road.

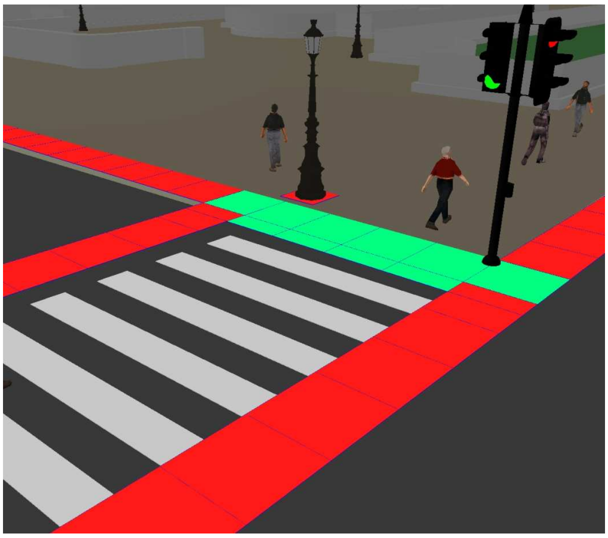

Fig. 3 - Border cells (red) prevent pedestrians to massively walk outside the appropriate areas.

2.1.4 Slowing down and stop

Each agent computes an *obstruction rate,* depending on the filling of the *direction table* (number of non-empty cells and average distance). A rate equal to zero means an empty table (no obstacle). Agents slow down if the rate becomes too high, but it will never cause them to stop. The number of non-empty cells has to be taken in account in order to avoid very close but small objects.

$$\tau = \sum_{i=0}^{N} \left(d_{\max} - d_i \right) \frac{N^P}{N}$$

Where N is the total number of cells, N^P the number of non-empty cells and d_{max} the maximal distance that it was initialized with.

An agent also slows down when someone walks too close in front of him, approximately at same speed and with the same orientation. Thus, they maintain a personal free space.

Agents are able to perceive traffic lights. If it is red for pedestrians, agents willing to cross the street will stop when they arrive at the border of the sidewalk or when they get too close to someone else waiting for the light to turn green (figure 4).

Fig. 4 - Agents waiting at a crosswalk.

2.2 Heterogeneity

In real life, crowds are very heterogeneous, both in terms of behavior and of appearance. This diversity is difficult to simulate but is a key to realism. We focused more on the behavior than on the visual aspect. From our observations, we identified three movement strategies: slow strategy, classical strategy and fast strategy.

Slow strategy: People walking slowly are either older persons or people going for a stroll. As they are about 50% slower than classical pedestrians, they don't care about distant obstacles, they only give attention to what is close to them. Their direction table is initialized with a small maximal distance.

Classical strategy: The majority of pedestrians follow this strategy. Classical pedestrians present an average behavior: they stay in appropriate areas, they slow down when too many people are in front of them, but they overtake if someone is too slow.
Fast strategy: Pedestrians that are rushing try to always walk at their maximum speed. They move about 50% faster than classical agents. Their obstruction threshold is higher, therefore they slow down less often than classical pedestrians. They are reckless: they don't give attention to border cells so they easily walk on the road if it allows them to overtake a pedestrian or to take a shorter path (figure 5).

The repartition of these strategies is important in order to get a realistic simulation: a majority of pedestrians must follow the classical strategy. A crowd composed of 80% classical, 10% slow and 10% fast pedestrians gave good results. These strategies bring heterogeneity and singular behaviors to the simulation. Fast agents do not respect the usual rules, like some people in real life.

Fig. 5 - A fast pedestrian overtakes slower agents and walks recklessly on the road.

2.3 Small Groups

In real life, we observe that more than half the people walk in small groups of two to six pedestrians. We counted on our videos 726 pedestrians, 43% of them walk alone, while 32% walk in pairs, 18% in groups of three people, 7% in groups of four people and the last 2% in groups of five or six people.
In our model, groups are composed of a leader and of followers. The leader decides of the speed and the direction of the group, followers copy their behavior on him. They all share the same targets. How to combine obstacle avoidance and group cohesion is an open question. For now, a group does not perform obstacle avoidance with moving objects. Alone pedestrians perceive groups as a single obstacle, they try to not cut through it (figure 6 and 7).

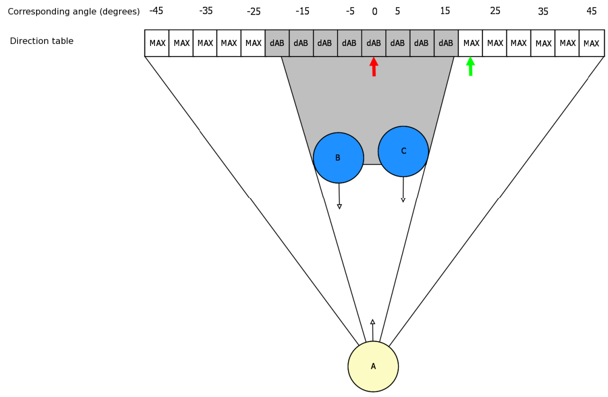

Fig. 6 - Agents B and C are part of a group, agent A sees them as a single obstacle, he will not walk between them.

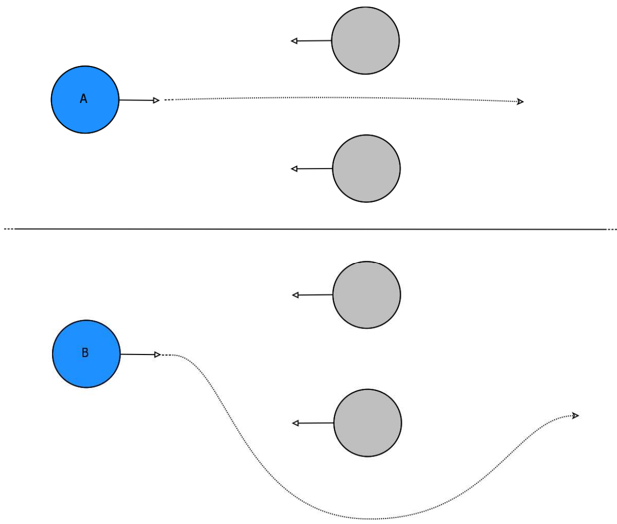

Figs .7 - Up, grey agents are not part of a group, A cuts between them. Down, grey agents are part of a group, B avoids them.

Each agent stores in its memory a list of other agents he knows. If an agent who walks alone (or is the leader of a group) meets one of them during simulation, they both will stop, stand a few seconds face to face and finally form a group. The leader of the new group is chosen arbitrarily. Fast pedestrians never stop when they meet a friend and do not form groups.

3. Results

Pedestrians smoothly avoid static and moving obstacles, stay in appropriate areas, have different behaviors and are able to form small groups. The main Ennis criteria are respected which brings great realism to the simulation. The different strategies allow singular perturbations like pedestrians crossing the street when and where they should not. We observe emergent behaviors like lanes formation in opposite flows (figure 8). Figure 9 compares a real scene with a simulated one.

The *direction table* technique is intuitive and can easily be implemented in any model allowing obstacle detection, with no need for any sophisticated environment. On a computer with an Intel Core 2 Duo 2,4GHz and 2 GB RAM our simulation can handle up to fifty agents without any lag at 30 frames per second (including 3D rendering). With 75 agents the framerate is down to 20 fps, we manage to keep an interactive rate (10 fps) with 200 agents.

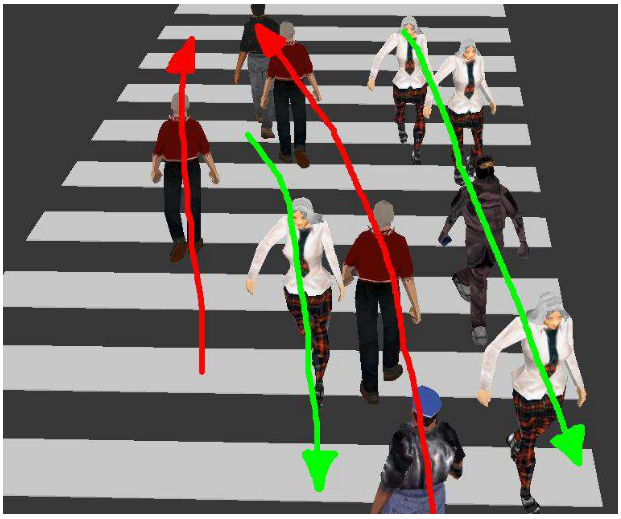

Fig. 8 - Lanes formation in opposite flow

Fig. 9 - Left, a real life photograph from downtown Toulouse. Right, a simulated scene with similar conditions.

In order to evaluate the pertinence of our choices, we ran a series of tests. Each test was based on the same principle: two simulations were launched initialized the same way, but one of them had a deactivated feature.

The first test concerned collision avoidance. We ran two simulations with fifty pedestrians initialized the same way (same positions, same strategies repartition, no small group), but agents of one of them did not perform collision avoidance (figure 10). It was visually obvious that having no obstacle avoidance ruins the realism. We also followed five agents on each simulation, during one minute, and counted how many times a collision occurred with one of them. 25 collisions occurred when avoidance is off, only 2 occurred when it is activated.

Fig. 10 - Left, collision avoidance is deactivated. Right, collision avoidance is activated.

We tested the visual impact of retaining pedestrians in appropriate areas with two simulations. Both were initialized with the same pedestrians at the same positions, one of the simulations had *border cells* but the other had not. The result is shown by figure 11, it is obvious that realism is enhanced when agents walk where they are supposed to.

Fig. 11 - Left, with no border cell, agents don't stay on sidewalks and crosswalks. Right, agents mainly stay on safe areas, it is obviously more realistic.

The impact of the presence of small groups was evaluated by the comparison between a simulation where every pedestrian is alone and another where some of them walk together, in small groups. The simulation with small groups seemed more natural (figure 12).

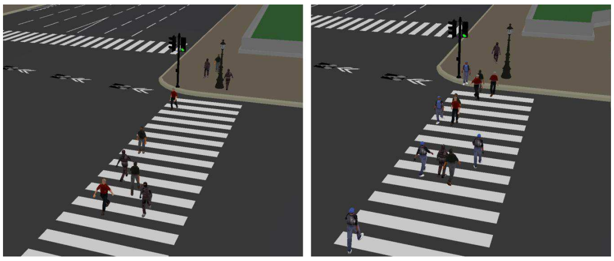

Fig. 12- Left, a simulation where everybody walks alone. Right, some agents walk in pairs or in groups of three persons.

To measure the effect of strategies heterogeneity, we ran a simulation where every pedestrian follows the classical strategy and another where the 80-10-10 repartition was respected. The first one gave the impression of a “clone army” since everyone was moving at the same pace. It is shown by figure 13. To enhance visualization, agents have been colored in accordance with their initial position. Left, a series of screen shots shows the evolution of a simulation where every agent follows the classical strategy. The crowd seems “frozen”. Right, 80% of the agents follows the classical strategy, 10% the slow strategy and 10% the fast strategy. There is more “mixing” between pedestrians, as fast agents overtake slow agents.

Strategies heterogeneity enriches the simulation, bringing singular behavior (very slow pedestrians, agents walking on the road, etc).

4. Conclusion

In this paper, we have presented our works based on an intuitive approach for pedestrians. The main characteristics of our system are to manage heterogeneous behaviours and cohesion of small groups. We have then substantially increased the realism of the simulations compared to existing methods.

This is the first step for managing a more complex environment where the characters can interact with the objects of the scene.

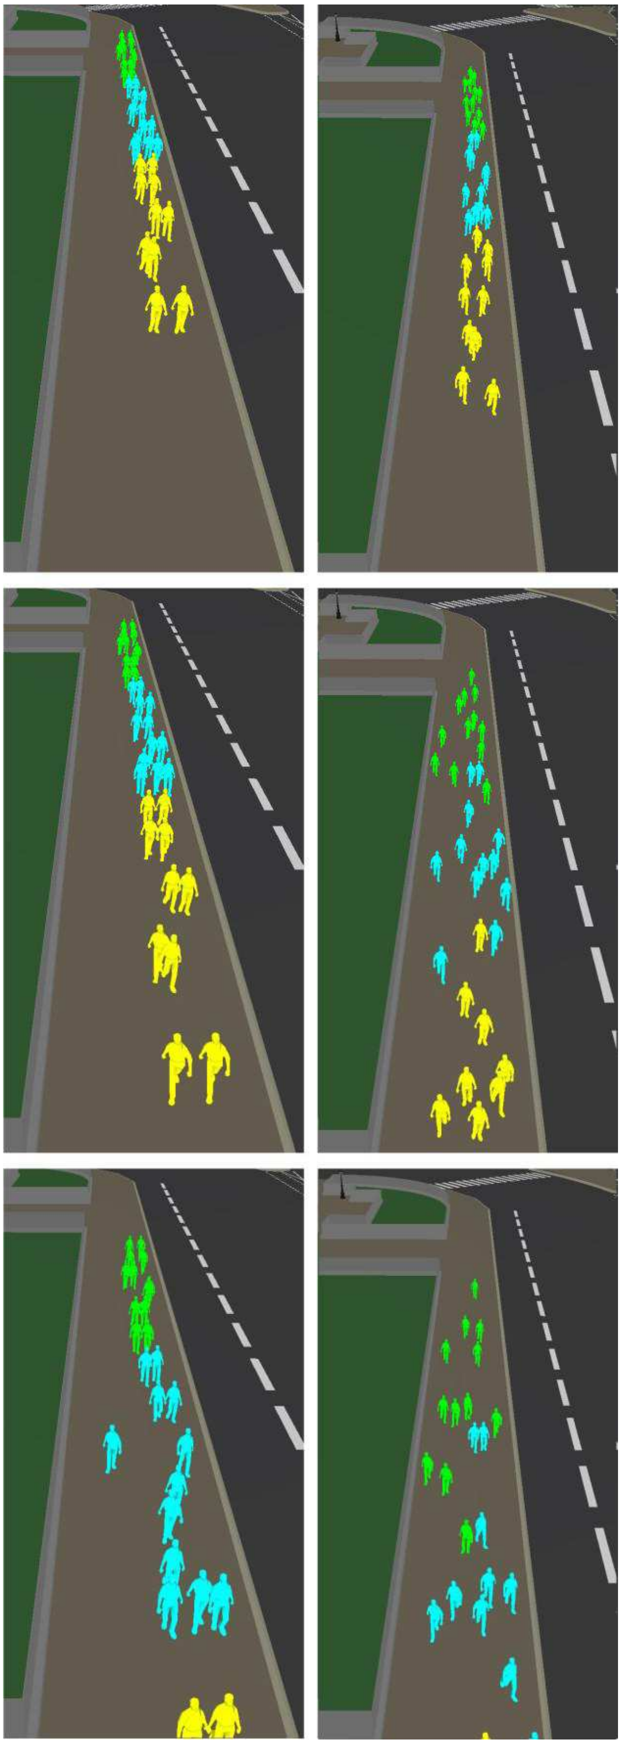

Fig. 13 - Left, classical strategy. Right, heterogeneous strategy.

The number of simulated characters remains the main problem. Further works focus on the optimization of the direction table. Each agent computes its own table while a table could be easily transmitted to a character moving in the same direction with few updates.

References

[Ennis10] C. Ennis, A. Gerdelan, C. O'Sullivan. *Plausible Methods For Populating Virtual Scenes.* CASA, 2010.

[Hughes03] R.L. Hughes. *The Flow of Human Crowds.* Annual Review of Fluid Mechanics, vol. 35, 2003.

[Lerner07] A. Lerner, Y. Chrysanthou, D. Lischinsky. *Crowds by Example.* Eurographics, 2007.

[Narain09] R. Narain, A. Golas, S. Curtis, M. C. Lin. *Aggregate Dynamics for Dense Crowd Simulation.* SIGGRAPH, 2009.

[Pelechano07] N. Pelechano, J.M. Allbeck, N.I. Badler. *Controlling Individual Agents in High-Density Crowd Simulation.* SIGGRAPH Symposium on Computer Animation, 2007.

[Shao05] W. Shao, D. Terzopoulos. *Autonomous Pedestrians.* Eurographics/ACM SIGGRAPH Symposium on Computer Animation, 2005.

[Treuille06] A. Treuille, S. Cooper, Z. Popovic. *Continuum Crowds.* SIGGRAPH, 2006.

[Yersin09] B. Yersin, J. Maïm, J. Pettré, D. Thalmann. *Crowd Patches: Populating Large-Scale Virtual Environments for Real-Time Applications.* Proceedings of Symposium on Interactive 3D Graphics and Games, 2009.