

Energy-efficient mid-term strategies for collision avoidance in crowd simulation

Julien Bruneau*, Julien Pettré†




Figure 1: In some situations, humans can see some kind of tunnel inside a flow that can facilitate the crossing of the flow. The left figure shows a first person view point where the tunnel is clearly visible. The right figure shows a top view of the situation as well as the space of colliding and collision free velocities. Regular local avoidance models try to find a collision free speed toward the goal. This means the red agent will choose a velocity from the green part of the space of velocities, making him slow down and miss the tunnel.

Abstract

When navigating in crowds, humans are able to move efficiently between people. They look ahead to know which path would reduce the complexity of their interactions with others. Current navigation systems for virtual agents consider the long-term planning to find a path in the static environment and the short term reaction to avoid collision with close obstacles. Recently some mid-term considerations have been added to avoid high density areas. However, there is no mid-term planning among static and dynamic obstacles that would enable the agent to look ahead and avoid difficult paths or find easy ones as human do. In this paper we present a system for such mid-term planning. This system is added to the navigation process between the path finding and the local avoidance to improve the navigation of virtual agents. We show the capacities of such system on several case studies. Finally we use an energy criterion to compare trajectories computed with and without the mid-term planning.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation; I.6.5 [Simulation and Modeling]: Types of Simulation—Animation;

Keywords: crowd dynamics, collision avoidance, interaction planning, navigation

1 Introduction

Crowd simulation is a very active field with many applications, like the entertainment industry in the aim of populating scenes for games and movies. For these applications we need to create virtual humans (agents) that behave in a believable manner. Spectators easily detect any weird behaviors, such as unnatural navigation trajectories. The main objective of this paper is to improve agents nav-

igation skills. In particular, we reinforce the capacity of the agents to setup relevant mid-term strategies for navigation among moving obstacles (typically, a crowd of moving agents).

A typical problem in agent navigation which remains quite neglected is how agents combine successive interactions in time. Today, most crowd simulators handle navigation in two steps. First, the long term strategy: the agent searches for a path between its position and its goal in the environment. This step is done by path finding modules. The second step is a short term strategy: while following the path the agent adapts its velocity to avoid any collision with nearby obstacles. This step is done by local avoidance modules. Most advanced solutions are able to combine simultaneous interactions: agents will find velocities that avoid collisions with all its neighbors obstacles at the same time. But such velocities do not always exist, especially when the agent is surrounded by many obstacles. In such a situation, local avoidance models have to reduce their constraints and allow colliding speed. Moreover the consequences of one local avoidance strategy (i.e., nature of the adaptations made) on the following interactions is generally not considered. Our objective is to prevent agents from performing an unnatural sequence of successive avoidance motions which can lead them to get stuck, whereas there are some more relevant and obvious solutions.

Real humans are very efficient to snake in and out between individuals or groups when crossing a flow or other dense crowds. An example is shown in Figure 1, where a red agent needs to cross a dense flow. There is an empty space in this flow that we call a tunnel. If we look at the first person view point on the left of the figure, the tunnel is clearly visible. A real human would probably detect it and use it to cross the flow with a minimum effort. Path finding techniques are unable to detect this tunnel as they plan inside the static environment and do not have information about the dynamic obstacles such as other agents. Local avoidance systems are also unable to detect this tunnel because they consider short term strategies only, they try to avoid any collision with only one velocity adaptation. The collision free velocity space and the colliding velocity space, shown on the figure, have been computed for a near future as done in local avoidance model. A closer look at the collision free velocities (in green) shows that the agent can either turn right and follow a parallel path to the flow for ever or turn left/reduce speed and pass behind the flow. But there is no indication of a tunnel in the flow on the different velocity spaces used by local avoidance to take a decision. In this situation local avoidance models make the red agent go straight forward and slow down a

*e-mail:julien.bruneau@irisa.fr

†e-mail:julien.petre@inira.fr

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org. SCA '15, August 7 - 9, 2015, Los Angeles, California. © 2015 ACM. ISBN 978-1-4503-3496-9/15/08...\$15.00 DOI: <http://dx.doi.org/10.1145/2786784.2786804>

lot when approaching the flow while the blue agents (in the flow) will move to let the red agent pass (see Figure 8). More recently, some heuristics have been proposed to choose a strategy depending on density, but these heuristics would also fail in our example as the tunnel is not straight and it requires considering several velocity adaptations to cross the flow. In this paper, we explore a third step for navigation in crowd simulators, which looks further into the future than the local avoidance step does, but not as far as the path finder does. Such a step is thus between path finding and local avoidance. It provides a mid-term strategy, allowing to handle the kind of situations shown in Figure 1.

There exist an infinite number of solutions to avoid a collision with an obstacle. An avoidance strategy can be defined as the nature of adaptations made to the trajectory to perform avoidance. A collision can for example be avoided by adapting speed or orientation, or a combination of both. The strategy choice has a huge impact on the resulting trajectories since, besides preventing the current collision, it also influences the relative motion with all the other obstacles, and thus changes the way the next potential collisions will be avoided. In this paper, we propose a solution to handle mid-term strategies by building several strategies to navigate through a crowd. These strategies consist of a list of adaptation sequences over time. The less costly strategy is then chosen. The cost of a strategy is a composition of the energy needed to follow the strategy, as well as an extra cost for every collision cause by the strategy. This allows agents to plan their way through a crowd, using empty spaces to navigate while optimizing their effort. We show the improvement on energy consumption compared to local avoidance systems.

Our contribution is a mid-term motion planning technique for complex scenarios with multiple moving obstacles. Our system, called Effective Avoidance Combination Strategy (EACS), is able to compute an energy-efficient avoidance path made of successive adaptations. To this end, it explores several possible ways to combine interactions, evaluates the energetic cost of each possible solution, and selects the most efficient one. The several options for avoidance paths result from the application of various avoidance strategies.

The remainder of this paper is organized as follows. In Section 2, we discuss the previous models used in agents navigation. In Section 3.1, we present our system and explain our approach. Section 4 shows the capabilities of our system on some specific situations, and an evaluation in terms of energy gain with respect to previous models. Finally, in Section 5, we draw some conclusions.

2 Related Work

In a crowd simulation, the navigation of agents generally results from two interlinked components. A global path planner is in charge of providing a sequence of way-points to guide agents to long-term goals. Some planners were specifically designed to handle crowds ([Pettré et al. 2006], [Bayazit et al. 2002], [Karamouzas et al. 2009], [Shao and Terzopoulos 2005]). While agents follow the planned motion, the local avoidance system avoids collision with moving obstacles that were not considered by the global planner ([Reynolds 1987], [van den Berg et al. 2008], [Helbing et al. 2000]). Between these two levels of motion synthesis, a relatively important gap exists.

Global path planners were extended to consider the presence of dynamic obstacles: a graph can vary through time to be adapted to dynamic obstacles [Sud et al. 2007] or information such as local density can be added to the graph [Pettré et al. 2006], including a prediction of its evolution over time [Kapadia et al. 2013]. However, these solutions are not able to estimate an optimal strategy over a sequence of interactions. Sud and colleagues [Sud et al. 2007] consider agents position only and do not anticipate future

agents' trajectory, thus preventing from exploring relevant strategies for successive interactions in time. Kapadia and colleagues [Kapadia et al. 2013] consider the case of multiple interactions by merging them to penalize going through denser areas and add time consideration to compute collision free velocities with neighbors in a similar way as velocity based local avoidance model. But they are not able to explore avoidance strategies for successive interactions. Our system could actually be used to extend this planning framework with a finer grain interaction analysis.

Local avoidance systems were also extended to consider some elements of mid-term strategies. Guy and colleagues [Guy et al. 2010] estimate the cost of applying a given strategy to perform collision avoidance. However, its consequence on the following interactions is not explored. Golas in [Golas et al. 2014] looks for the density of the crowd further ahead and penalizes avoiding velocities that go toward high density area. Using the same principle, [Best et al. 2014] adapt the walking speed to the crowd density ahead of the agent. These methods allow nice improvements over agent's navigation but still limit themselves to a single velocity adaptation which is not sufficient to solve situations such as the one presented in the introduction and shown in Figure 1. The method presented in [Kapadia et al. 2009] use egocentric affordance fields to find a path through dynamics and static obstacles and should be able to solve the situation of Figure 1. While the egocentric field should contain enough information on the immediate surroundings of the agent to find such narrow passageways, information are lost over distance and narrow passageways won't be detected if too far away even if there is no obstruction between the agent and the passageways.

Our approach differs from the previous work in various aspects. The proposed method is based on an accurate evaluation of the cost of several feasible paths, it is not guided by heuristics which only consider the effect of density on navigation efficiency. Our method finds good strategies to combine interactions in time by selecting strategies that may improve several successive avoidance interactions. It differs from techniques which apply a greedy approach by selecting the best strategy for the next interaction only, such as [Best et al. 2014]. The most similar approach to ours is probably the one by Guy and colleagues [Guy et al. 2010], with the major difference that we explore the n next interactions occurring in time, whereas this previous work only considers the on-going interactions with neighbor agents.

3 Effective Avoidance Combination Strategy

3.1 Overview

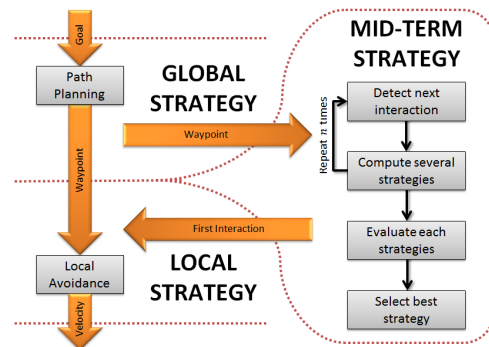


Figure 2: Navigation system architecture with three different levels of decision.

The objective of the EACS system (EACS stands for Effective

Avoidance Combination Strategy) is to enable agents to perform efficient navigation through other moving agents.

The typical structure of a crowd simulation framework consists of two levels. As presented in Figure 2, there is the path planning for global strategy and the local avoidance for local strategy. The path planning find a path in the environment which consists of a sequence of positions without any fixed obstacles between two consecutive positions. The local avoidance uses the first position of the path, which is reachable in a straight line from the agent's position, and computes a speed to reach it without any collision with nearby fixed or moving obstacles. The EACS system stand between these two modules, it takes the reachable first position of the path from the path planning as a mid-term goal and output a preferred velocity to the local avoidance system. The local avoidance system will then follow this preferred velocity while avoiding very close obstacles. Efficient navigation is obtained by setting the relevant avoidance strategy for the few next future interactions with anticipation.

To this end, EACS detects all next interactions on the path to the mid-term goal and computes a sequence of velocity adaptations to successively perform collision avoidance. The variety of possible adaptations to achieve collision-free motion results into different mid-term navigation paths. EACS evaluates and compares their efficiency from the energetic point of view and selects the most efficient one to steer agents. In the remainder of this section we start by defining the main concepts used in our system, followed by a description of our technique.

3.2 Definitions

Collision Course: two agents are on a collision course when their position and velocity vectors are so that their future distance of closest approach is below contact distance if velocity vector is kept constant.

Avoidance Strategy is the nature of the adaptations made to velocity vectors to avoid future collisions. For example, an agent can avoid a collision with another agent by adapting speed (i.e., the norm of the velocity vector) or by turning (i.e., by rotating the velocity vector).

Interaction Waypoints (IW): for two agents on a collision course and performing collision avoidance, the IW are the space-time waypoints by which agents go when avoiding.

Cardinal Interaction Waypoints (CIW) are the 4 IW which result from 4 specific avoidance strategies with the minimum required amount of adaptations: accelerating only, decelerating only, turning left only and turning right only. We note these cardinal interaction points by \mathbf{a} , \mathbf{d} , \mathbf{l} and \mathbf{r} , respectively.

Interaction Paths (IP) are a sequence of IW ordered by increasing time. An agent that follows an IP will go through all its IW. Between two consecutive IW the agent follows the straight line between the two positions at a constant speed.

Interaction Segments (IS) are a part of an IP composed of two consecutive IW.

Goals (G) are special IW that do not have a time constraint. A goal is a position in space that agents try to reach, regardless of the when they reach it. Goals appear only at the end of an IP or IS. When an agent is following an IS between a regular IW to a goal, it simply walks at its preferred speed toward the goal.

3.3 Cardinal Interaction Points (CIW)

In this section, we describe how we compute the 4 CIW. We arbitrarily chose this four points, they cover the four types of adaptation

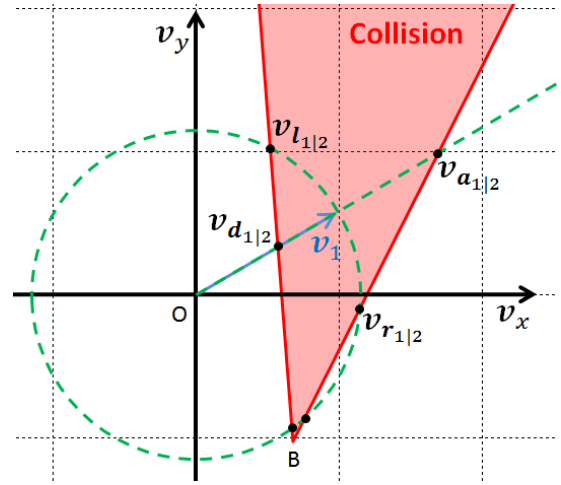


Figure 3: Collision velocity space for one interaction, in red. $\mathbf{v}_{\mathbf{a}_{1|2}}$, $\mathbf{v}_{\mathbf{r}_{1|2}}$, $\mathbf{v}_{\mathbf{d}_{1|2}}$ and $\mathbf{v}_{\mathbf{l}_{1|2}}$ represent the collision free velocities that are used to compute the CIW.

(accelerating, decelerating, turning right and turning left) as well as the two orders of passage (passing before the other agent or after). Using these 4 CIW yields satisfying results but other IW could be consider when computing avoidance strategies.

The 4 CIW computation is described for a given agent α_1 on a collision course with agent α_2 . An example is presented in Figure 4. Their respective 2-dimensional positions and velocities at current time t_0 are noted $(\mathbf{p}_1, \mathbf{v}_1)$ and $(\mathbf{p}_2, \mathbf{v}_2)$. The CIW (3-dimensional: space and time) are noted: $\mathbf{a}_{1|2}$, $\mathbf{d}_{1|2}$, $\mathbf{l}_{1|2}$ and $\mathbf{r}_{1|2}$. To compute these points, we base our principle on the velocity-obstacle introduced in [Fiorini and Shillert 1998] as illustrated in Figure 3. The key idea is to compute the 4 admissible velocities which enable collision free trajectories and which correspond to 4 specific strategies (deceleration, acceleration, turning left, turning right). We recall that the set of admissible velocities for α_1 concerning its interaction with α_2 is noted $AV_{1|2}$ and is defined as follows:

$$AV_{1|2} = \{\mathbf{v}_1 \in V_1 \mid \forall t \in [t_0, t_0 + \tau], \text{dist}_{1,2}(\mathbf{v}_1, t) \geq c\} \quad (1)$$

where: V_1 is the set of all the reachable velocities for α_1 , τ is the size of a time window, $\text{dist}_{1,2}(\mathbf{v}, t) = \|(\mathbf{p}_2 + \mathbf{v}_2 t) - (\mathbf{p}_1 + \mathbf{v} t)\|$ and c is the collision distance threshold.

Then, the CIW define trajectories followed by α_1 with specific strategies (velocities): $\mathbf{v}_{\mathbf{a}_{1|2}}, \mathbf{v}_{\mathbf{r}_{1|2}}, \mathbf{v}_{\mathbf{l}_{1|2}}, \mathbf{v}_{\mathbf{d}_{1|2}} \in AV_{1|2}$:

- $\mathbf{v}_1 \cdot \mathbf{v}_{\mathbf{a}_{1|2}} = 0$, $\|\mathbf{v}_1\| < \|\mathbf{v}_{\mathbf{a}_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{\mathbf{a}_{1|2}}, t) = c$,
- $\mathbf{v}_1 \cdot \mathbf{v}_{\mathbf{d}_{1|2}} = 0$, $\|\mathbf{v}_1\| > \|\mathbf{v}_{\mathbf{d}_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{\mathbf{d}_{1|2}}, t) = c$,
- $\det(\mathbf{v}_1, \mathbf{v}_{\mathbf{l}_{1|2}}) > 0$, $\|\mathbf{v}_1\| = \|\mathbf{v}_{\mathbf{l}_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{\mathbf{l}_{1|2}}, t) = c$,
- $\det(\mathbf{v}_1, \mathbf{v}_{\mathbf{r}_{1|2}}) < 0$, $\|\mathbf{v}_1\| = \|\mathbf{v}_{\mathbf{r}_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{\mathbf{r}_{1|2}}, t) = c$.

The CIW $\mathbf{a}_{1|2}$, $\mathbf{d}_{1|2}$, $\mathbf{l}_{1|2}$ and $\mathbf{r}_{1|2}$ are then:

- $\mathbf{a}_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{\mathbf{a}_{1|2}} \cdot \text{time}(\mathbf{v}_{\mathbf{a}_{1|2}}), \text{time}(\mathbf{v}_{\mathbf{a}_{1|2}}))$,
- $\mathbf{d}_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{\mathbf{d}_{1|2}} \cdot \text{time}(\mathbf{v}_{\mathbf{d}_{1|2}}), \text{time}(\mathbf{v}_{\mathbf{d}_{1|2}}))$,

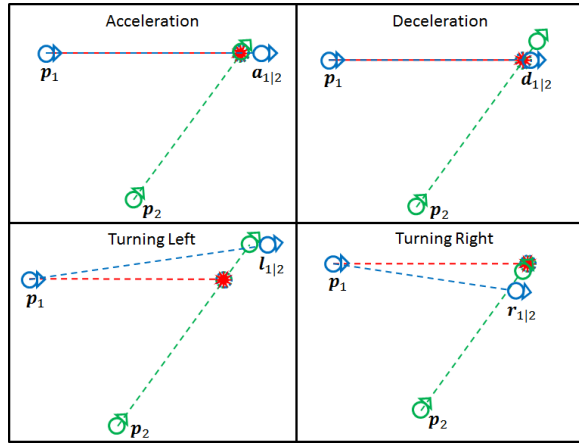


Figure 4: Example of collision solved using the four CIW. The collision trajectories are drawn in red and the four trajectories followed to reach the CIW are shown in blue. The position of each CIW are shown together with the position of the other agent when the CIW are reached.

- $l_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{1|2} \cdot \text{time}(\mathbf{v}_{1|2}), \text{time}(\mathbf{v}_{1|2}))$,
- $r_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{r_{1|2}} \cdot \text{time}(\mathbf{v}_{r_{1|2}}), \text{time}(\mathbf{v}_{r_{1|2}}))$,

where $\text{time}(\mathbf{v})$ is the time of closest approach ($ttca$) delayed by a small amount: $\text{time}(\mathbf{v}) = ttca(\mathbf{v}) + \min\left(\frac{c}{v}, \frac{c}{v_2}\right)$; $ttca(\mathbf{v})$ is computed assuming agents α_1, α_2 travel with velocities \mathbf{v}, \mathbf{v}_2 .

This delay allows the two agents to get away from each other to avoid resetting them on a collision course if one of the two agents performs a new maneuver too soon (e.g., to solve one next interaction). The example presented in Figure 4 clearly shows that once the CIW is reached, the collision risk with the other agent is avoided.

3.4 Collision Free Interaction Path

The Collision Free Interaction Path construction is presented on Algorithm 1 and illustrated with an example in Figure 5. The collision free IP is built iteratively by the algorithm. It starts with the straightforward IP to the goal (line 3). A collision test is performed with this IP (line 7). If no collision is found, a collision free IP is detected (line 18). If a collision is found, the IP is discarded and the four CIW are computed and used to build four new IP (line 8). The test is then repeated on these new IP, first on the IS right before the interaction (line 9 to 12) then the IS from the interaction to the goal (line 16) until one IP reaches the goal (line 18).

An example is presented in Figure 5. The initial situation is shown in Figure 5A): the blue agent starts at the position s and has to reach the position g . As we said the algorithm starts by testing the straight forward IP (Figure 5B)), a collision is detected and four new IP are created $sl_{1|3}g, sd_{1|3}g, sa_{1|3}g$ and $sr_{1|3}g$. Now, these four new IP have to be tested starting with the IS just before the interaction: $sl_{1|3}, sd_{1|3}, sa_{1|3}$ and $sr_{1|3}$. Figure 5C) shows the test for the turning left solution $sl_{1|3}$. A collision is detected on the IS and again four new paths are created. Figure 5D) shows the test for the accelerating solution $sa_{1|3}$ and no collision is found. But this is only the first part of the IP, there is still the IS $a_{1|3}g$ to test. This test is done in Figure 5E). A collision is detected and only two paths are created. Both agents are going toward each other, they cannot avoid the collision by accelerating. Decelerating will not solve the

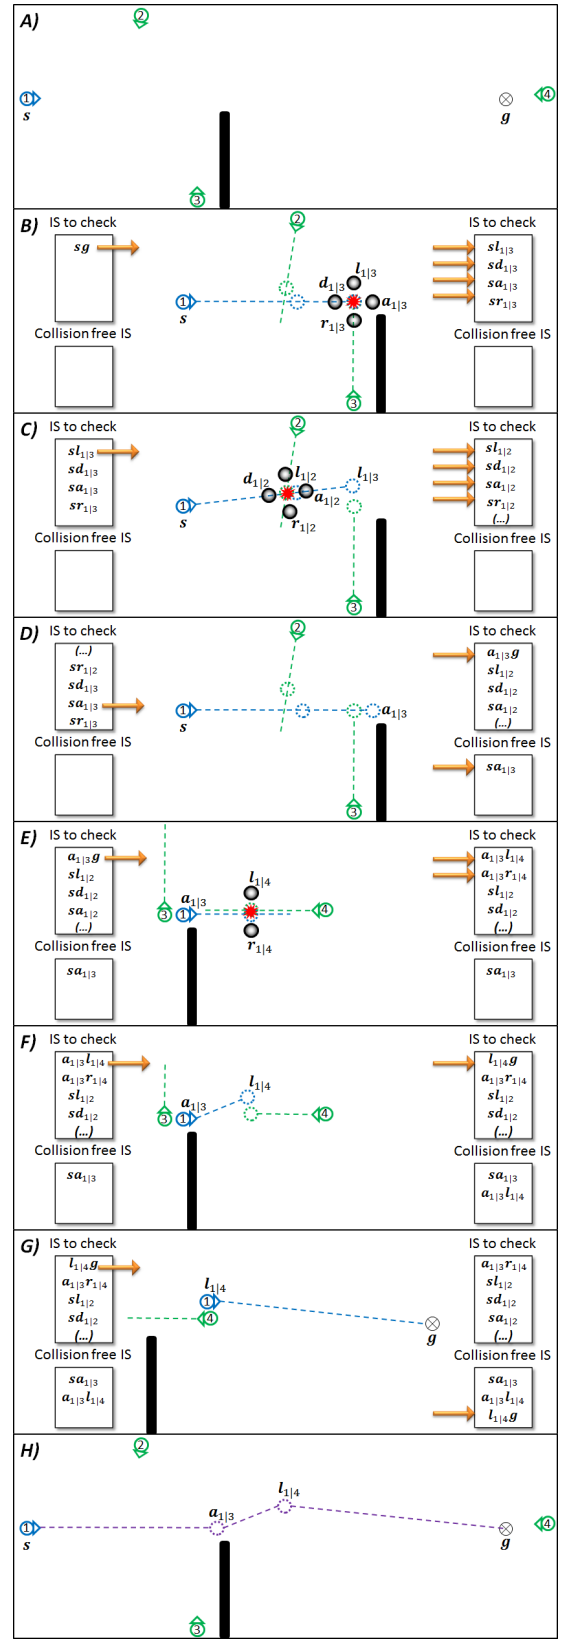


Figure 5: Example showing the construction of a collision free IP illustrating the different steps taken by the EACS system.

Data: Starting position s , Goal g

Result: Build a collision free IP

```

1 checkList←[];
2 collisionFreeList←[];
3 add(sg, checkList);
4 pathFound←FALSE;
5 while !pathFound do
6   bc←pullSegment(checkList);
7   if collision(bc) then
8     a,d,l,r ← computeCIW(bc);
9     add(ba,checkList);
10    add(bd,checkList);
11    add(bl,checkList);
12    add(br,checkList);
13  else
14    add(bc, collisionFreeList);
15    if c!=g then
16      add(cg, checkList);
17    else
18      pathFound←TRUE
19    end
20  end
21 end
22 P←buildPath(s,g,collisionFreeList);
23 return P;
```

Algorithm 1: Algorithm used by EACS to build collision free Interaction Paths

problem either but only delay it as the blue agent will have to go toward the goal eventually. The turning left solution is tested on both IS, first on $a_{1|3}l_{1|4}$ (Figure 5F) then on $l_{1|4}g$ (Figure 5G)), and no collision is found. The goal is reached and a collision free IP is built: $sa_{1|3}l_{1|4}g$ (Figure 5H)).

3.5 Cost and ranking

On the example presented in Figure 5, the collision free IP $sa_{1|3}l_{1|4}g$ is built. But if the turning right solution for the last interaction had been tested before the turning left one, we would have ended with the collision free IP $sa_{1|3}r_{1|4}g$. With another order of test, we could have also built the collision free IP $sl_{1|2}g$. In the end, there are many different IP to go through a crowd. We are looking for only one but one that is efficient and would most likely be picked by a real walker. To that end, we used a cost function to rank the different IP and select a single one among them. We chose to use the energy consumption as the cost function. Real humans often favor the least energy consuming way to perform a task [Zipf 1949], moreover it has already been used to improve collision avoidance models [Guy et al. 2010].

The energy formula, $E = m \int (e_s + e_w |\mathbf{v}|^2) dt$, is from [Whittle 2003]. The values of $e_s = 2.23 J.s.Kg^{-1}.s^{-1}$ and $e_w = 1.26 J.s.Kg^{-1}.m^{-2}$ minimize the energy consumption per distance for a speed of $|\mathbf{v}| = 1.33 m.s^{-1}$ which will become the preferred speed of the agent. We want to have agents with different preferred speed. From the energy formula per distance, Equation (2), we set a constant distance to look for the speed that minimizes the energy consumption: \mathbf{v}_{pref} (see Equation (3)). Then we are able to change the value of e_w for the chosen \mathbf{v}_{pref} (see Equation (4)).

$$E = m \left(\frac{e_s}{|\mathbf{v}|} + e_w |\mathbf{v}| \right) d \quad (2)$$

$$\frac{dE}{dv} = m \left(\frac{e_s}{|\mathbf{v}|^2} + e_w \right) d = 0 \quad (3)$$

$$e_w = \frac{e_s}{|\mathbf{v}_{pref}|^2} \quad (4)$$

In some situations, allowing collisions is necessary to find good solutions (see Figure 10). So we enabled the search algorithm to consider IP with collisions. To keep a preference for collisions free IP, an extra cost is added for IP with collisions. This extra cost is the equivalent of making a long detour: the energy consumption at preferred speed for a specific distance which is a parameter of the system ($CollisionCost = m(\frac{e_s}{|\mathbf{v}_{pref}|} + e_w |\mathbf{v}_{pref}|) d_{collisionDetour}$). This way, agents will still go around dense area when these are not too big, and will go through them otherwise.

3.6 Following case

When there is someone in front going in the same direction but slower, there is not always enough space to avoid him. In this case, people can try to squeeze to avoid him despite the lack of space or they can match the speed of the one in front to follow him until there is enough space to overtake him or until their trajectories separate. The following behavior is close to collision avoidance as it can be used to prevent a collision. At the same time, it is very different as the collision is never solved by following, only postponed. As said before, people follow until they can properly avoid the one in front of them or the collision risk disappears by itself. When computing the four CIW, the agent follows a velocity that prevents the collision until the collision risk no longer exists. As explained above, when following the collision risk remains except if one of the agent changes direction. So the agent can still change his speed to match the one of the agent in front and prevent a collision by following but we need to determine how long should it follow to compute the corresponding IW.




Figure 6: Detection of a following case and the two Interaction Waypoints i and j computed from the following speed

Following happens when the tested IS has a collision with an agent that goes in the same direction, this situation is represented in Figure 6. Whether the IW h (the end of the IS) is the goal or just an intermediate IW place here to avoid another agent, the position and/or the time of the IW is important. This is why it has been decided that when following, two IW should be computed. For the first one j , the agent will follow until it reaches the previous position he wanted to reach (the position of h). For the second one i , the agent will follow for as long as it would have walk with the previous IS kh .

From the Figure 6, we have the velocity of the green agent \mathbf{v} , the IW $k = (p_k, t_k)$ and $h = (p_h, t_h)$. We then compute the IW $j = (p_h, \frac{p_h - p_k}{\mathbf{v}} + t_k)$ and $i = (\mathbf{v} * (t_h - t_k), t_h)$.

4 Results

In this section we present the results of our work. The simulations have been done by using our system coupled with RVO2. EACS computed a strategy and then set the preferred velocity required by RVO2. No additional path finding algorithms have been used as the studied situations were simple enough and the goal is always reachable from the agents' starting position.

4.1 Performance

The EACS system computes and evaluates many different IP. An estimation of the final cost of an IP is done to advance the construction of the most promising one in priority. Three limitations further bound the exploration to avoid an explosion of the computational cost: maximum planning time (we don't plan further than one minute ahead), maximum number of interactions per plan (we don't build strategies of more than 20 interactions) and a maximum number of main loop iterations (500).

Moreover, much like path planning, interaction planning does not need to be performed every step. Plans are not made to be discarded immediately afterwards, except in some exceptional situations. The agents thus compute their IP every ten steps (defaults time step is 0.1s so agents are updating their plan every seconds). In order to spread the computational cost of the interaction plan, the system does not compute the IP of every agent at once. Instead, it computes the IP of a tenth of the agents at every step.

These measures and the use of multiple threads during the planning enable the different simulations presented in this section to be run in real time. Moreover computation time has been measured for several counter flow simulations with different number of agents (an example of such simulation is presented in the companion video). The results, shown in Figure 7, have been measured on a computer with an Intel®Xeon®Processor E5-1603 (4 cores, 10M Cache, 2.80 GHz, 0.0 GT/s Intel®QPI) and 32GB Memory. In this situation where every agents have to compute complex IP, we are able to run simulations with up to 200 agents in real time with the default time step of 0.1s.

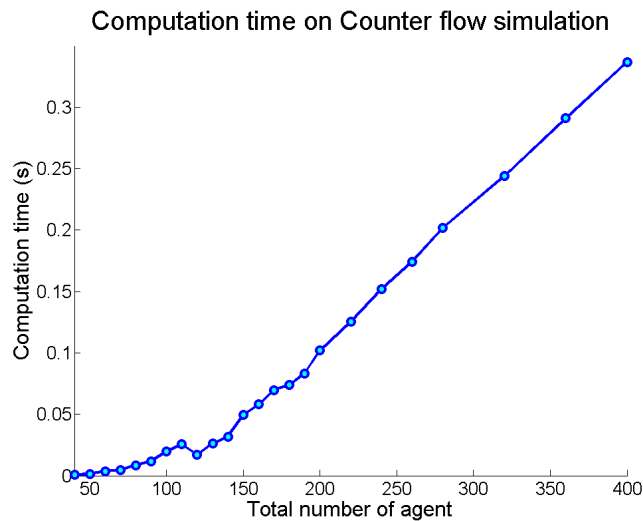


Figure 7: Computational time for one step of a counter flow simulation depending on the number of agent in the simulation.




Figure 8: Simulation of an agent (in red) going through a flow, with the flow having a tunnel to facilitate the crossing. Left: simulation including EACS; Right: simulation without EACS.

4.2 Case study

Several situations were designed to test the EACS system and check the improvement it brings to crowd simulation. The first one is the example used in the introduction to show the limitation of the current models. The results are shown in Figure 8. As expected, RVO2 alone is unable to use the tunnel to go through the flow. Instead, the red agent goes straight forward and has trouble going through the flow as it spends around 10 seconds in the flow itself. With the EACS system, the red agent is able to take the tunnel which facilitates its navigation greatly: the agent spends only 2 seconds in the flow.




Figure 9: Simulation of a clogged exit, EACS is able to choose other exits that are further away but not clogged.

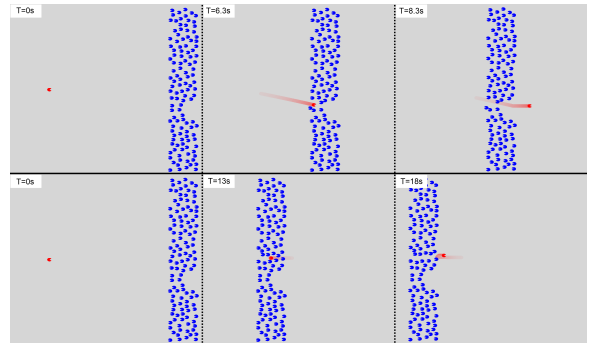


Figure 10: Simulation of an agent (in red) going through a flow that goes in the opposite direction. The flow has no tunnel but a section of it has a lower density. At the top, we have the simulation with EACS and at the bottom the simulation without EACS.

The second one is shown in Figure 9 and is very similar to the first one. Instead of going through a flow, agents need to exit or enter a building that has 3 entrances. All the agents are close to one exit which quickly becomes clogged. As for the first situation where the agent was able to detect the tunnel, in this situation, some of the agents detect that the other two exits are empty. Instead of waiting and avoiding many obstacles which can be painful, they decide to go through the two others exits. The extra cost from the detour is largely compensated by all the required velocity adaptations to go

through the clogged exit. This is easily witnessed in real life: in the train station for example, where people will use further exits to avoid the dense crowd of the closest one.

In the third situation shown in Figure 10 there is no collision-free solution to cross the incoming flow. Such situations exist in real life where we have to cross a very dense crowd and we require other people to give us some space to go through. In the presented situation, even if there are no collision-free paths through the crowd, there is a weak point where the density is lower. Going through this point would make it easier to go through the flow. EACS is able to detect such weak points in the flow and compute a path through it. If we compare with the trajectory computed without EACS (at the bottom), we can clearly see that it is easier to use the weak point to go through the flow. In the end, the agent controlled by RVO2 spends around 10 seconds in the flow while the one controlled by EACS spends only 2 seconds in it.

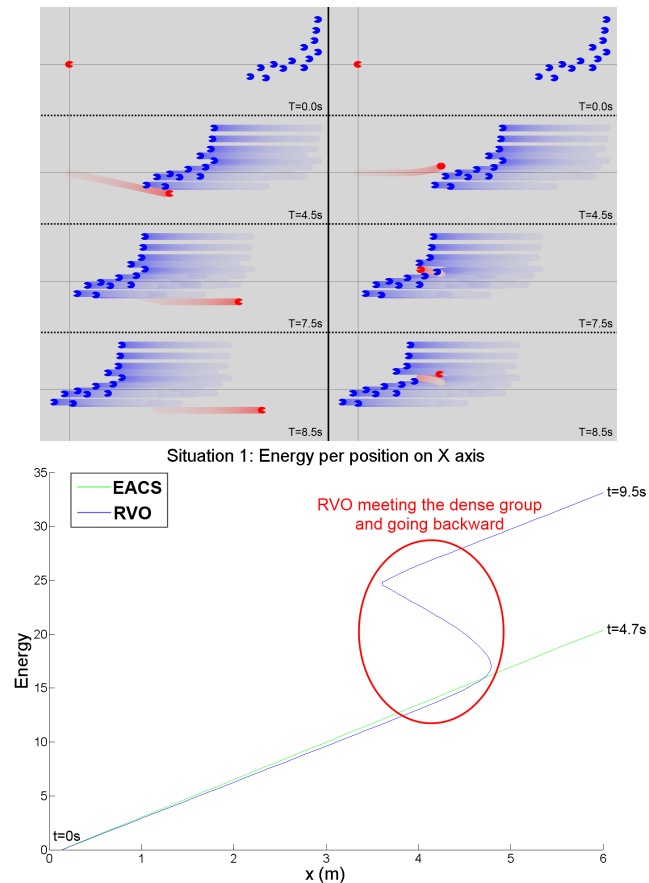


Figure 11: Top: simulated test where short-term cost leads to bad mid-term cost, agents simulated with EACS on the left and without on the right. At the bottom, the graph quantifies the energy consumption of the red agent (For RVO2 and EACS) as they walk towards the goal (In this situation, it is equivalent to the distance traveled on the x axis). The section of the blue curve circle in red correspond to the RVO2 agents going backwards to avoid the dense group (T=7.5s on the top figure).

Another situation is shown in Figure 11 and has been designed to show the difference in optimizing the short-term outcome and the mid-term one. If we look at the trajectory given by RVO2 alone (top right), we can see that the red agent is avoiding by going on the left. If we consider only the first few blue agents, it is the best solution as going on the right would require to go faster and to shift

more on the side. This is confirmed by the energy graph, the energy consumption of EACS which goes on the right is higher than the one of RVO2 for the first 3m. But when we look further, we can see that the agent will meet with more people when going to the left. This dense area is so hard to navigate that the agent will even move backwards. Looking at the graph, we can notice the sudden rise in energy consumption for RVO2 around 4m as well as the curve going backwards because the agent walks in the opposite direction of its goal. The curve of EACS energy consumption remains stable which confirms that the EACS strategy is better on the mid-term.

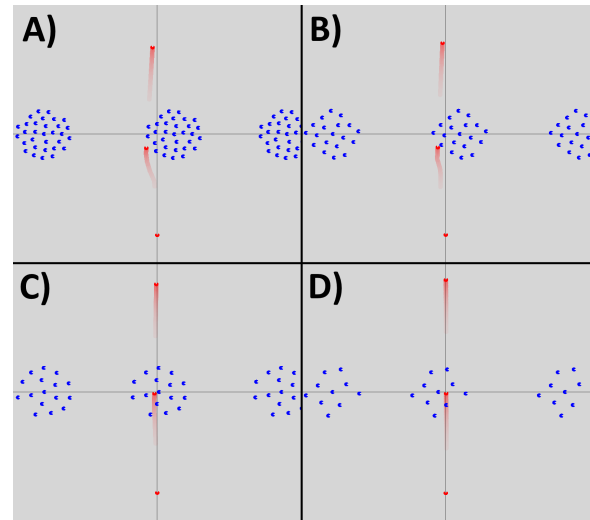


Figure 12: Simulation of group avoidance for different group densities.

The last situation is about group avoidance. Previous experiments have shown that real humans go around dense groups and through sparse ones [Bruneau et al. 2015]. We have made several simulations where an agent had to avoid groups of different densities to check that EACS is able to reproduce the real trend. We can see some of them in Figure 12, the situation setup is the same as the one in the experiment with real humans. The red agent is able to decide whether to go through the group or around it and reproduces the same trend as that of a real human: it circumvents the two dense groups (Figures 12A and 12B) and traverses the sparse ones (Figures 12C and 12D).

4.3 Energy

To check the energy gain of EACS over pure RVO2, we studied the trajectories of agents going through an orthogonal flow. We tried several different conditions to analyze their impact on the results. We performed 2800 simulations: 7 density types x 19 exploration limits for EACS and 1 without EACS x 20 repetitions. For each simulation, a flow was randomly generated and 10 agents were randomly placed on the same side of the flows with their goal on the other side. The energy consumption of the 10 agents were computed from their starting position to their goal location and then averaged.

The first condition studied is the density type. For each generated flow, the same dimensions (10mx80m) and number of agents (700) were used. Thus the global density was the same for each flow. The variation among the different conditions was made on the regularity of the density in the flow. These variations are shown in Figure 13. We have a uniform flow on the right and the different types of irregularity until the very irregular flow on the left. These irregularities

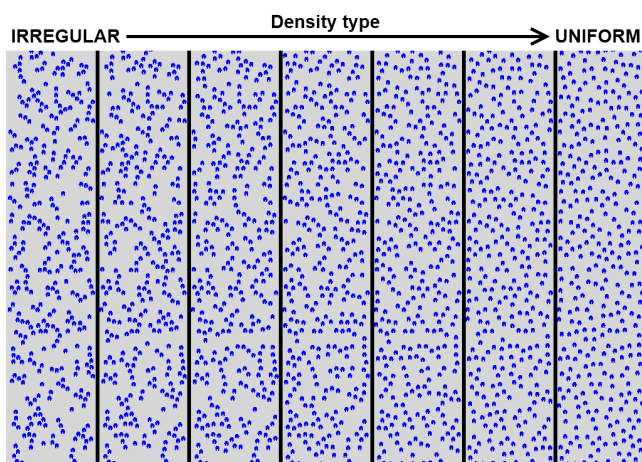


Figure 13: The different density types used for the energetic study from very irregular (left) to uniform (right).

create small dense clusters and leave some space in between to navigate more easily. We expect EACS to be able to use this in-between space and perform better when going through irregular flows.

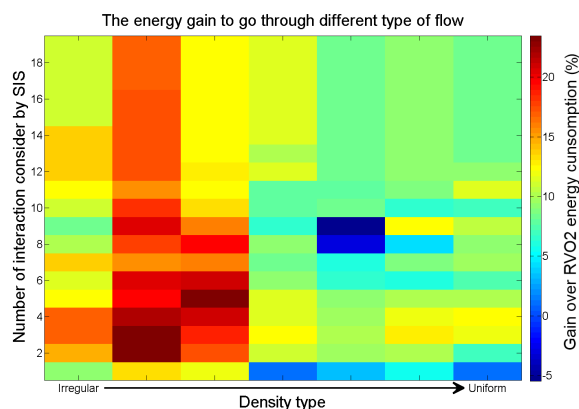


Figure 14: This graph shows the amount of energy preserved by the agents when using EACS to optimize their avoidance strategies compared to EACS-less strategies.

The second condition studied is the exploration limit for EACS. When EACS is looking for a path it will sequence several interactions over time. The number of interactions can become very big in some situations which can cause performance issues. Moreover it is doubtful that real humans can consider a high number of future interactions. For this reason, we set a maximum number of interactions when constructing an IP in EACS. When an IP reaches the limit, it will not be extended anymore. We expect that a low maximum number of interactions will lead to results close to RVO2. We also expect that a very high number of interactions will perform similarly or worse than a medium number of interactions: the further we plan in the future, the less the extrapolation is relevant.

For the results, for each exploration limit, we compare the average energy consumption with and without EACS. The Figure 14 shows this consumption difference in percent. As expected, EACS is performing better for very irregular densities. When the density becomes uniform, the energy consumption with and without EACS becomes similar. If we look at the different results for different maximum numbers of interactions, we can see that having a value

higher than 14 does not seem to change the results much. Indeed, optimum results seem to be found for a maximum number of interactions around 4.

4.4 Limitations

Our system has limitations. The main one is the number of future interactions explored by the system. We showed in the previous section that the system is particularly efficient when used in specific contexts, such as when agents density is not uniformly distributed in space. We can actually easily guess what type of crowd motions will enable EACS to be efficient. For example, we expected that non-uniform density distributions would open larger efficient interaction paths in the crowd with a great probability that these paths remain valid in spite of all agent adaptations.

More generally, the EACS system is particularly efficient when the initial predictions are good and turn out to be representative of the actual motion of neighbor agents. In the opposite case, predictions turn out to be false, and the energy gain by EACS is poor. We also observe in the energetic evaluation of EACS solutions (previous section) that considering interactions too far in the future decreases the quality of the solutions provided by EACS. Note that EACS guided agents toward fewer counterproductive situations, when compared with naive short-term solutions.

One solution to this could be to evaluate on-line the quality of the predictions made by EACS. The number of interactions explored in the future could be directly dependent on the quality of this prediction. If predictions are good, agents can explore wider future time windows and more numerous future interactions, and conversely.

Another limitation of EACS is to only consider an energy-related cost function. We are convinced that this criterion is considered by human walkers when setting mid-term strategies, but we are also convinced that other factors play a role. For example, we could check how close the solution paths to other agents motion are. Additional social distances could be considered in the cost of paths. As another example of a social criterion, passing in front of an agent at close distance could be penalized in comparison with the strategy of giving way, which could be considered to be more polite.

Finally, EACS explores a subset of all the possible strategies to perform collision avoidance. We make an arbitrary choice to explore adaptations exclusively made of speed or orientations changes, whereas all intermediary solutions could be explored. We are here in a typical trade-off between computation times and quality: we could add more than 4 CIW in the system to consider mixed adaptations strategies, but this would result into higher computation times.

5 Conclusion

We have modified the usual navigation process of crowd simulation to add some mid-term consideration. To this end we have designed an interaction planner, the Effective Avoidance Combination Strategy that creates interaction path which enables agents to consider non straight paths to go through a crowd. As results, we have shown with specific test cases that the agents are able to select strategies that use density tunnel to facilitate their navigation through a crowd. Moreover, while building the interaction path, the EACS evaluates the repercussion of its avoidance choices further in the future. This allows agents to make some extra effort in resolving their current collision so as to simplify future interactions.

With an energetic study, we show that by adding the EACS to the regular navigation process agents use less energy to navigate. The

study has shown that the EACS is especially useful to improve trajectories when facing crowd with irregular densities. The study has also highlighted that too much mid-term planning is not beneficial, which can be explained by the incapacity to correctly extrapolate the other agents' trajectories too far in the future. It seems that planning around 4 interactions ahead yields the best results.

We have shown the improvement in the quality of agents navigation brought by our work. The next step is to compare it with real human behaviors and interactions planning. Especially, it would be interesting to study how many interactions ahead real humans plan. Another interesting direction is to work on the evolution of the interaction path. Right now, the IP is computed at a constant interval and each time the computation starts from scratch. Real humans are known to be highly adaptable creatures, in most cases they most likely adapt their plan based on changes in the environment instead of planning a new one from scratch. This would greatly improve performance but it might also improve results by making the plan more reactive to the environment changes.

Acknowledgements

This work is funded by the French National Research Agency ANR, project PERCOLATION number ANR-13-JS02-0008.

References

- BAYAZIT, O. B., LIEN, J.-M., AND AMATO, N. M. 2002. Roadmap-based flocking for complex environments. In *Computer Graphics and Applications, 2002. Proceedings. 10th Pacific Conference on*, IEEE, 104–113.
- BERG, J., GUY, S., LIN, M., AND MANOCHA, D. 2011. Reciprocal n-body collision avoidance. In *Robotics Research*, C. Pradalier, R. Siegwart, and G. Hirzinger, Eds., vol. 70 of *Springer Tracts in Advanced Robotics*. Springer Berlin Heidelberg, 3–19.
- BEST, A., NARANG, S., CURTIS, S., AND MANOCHA, D. 2014. Densesense: Interactive crowd simulation using density-dependent filters. In *Symposium on Computer Animation*, 97–102.
- BRUNEAU, J., OLIVIER, A.-H., AND PETTRÉ, J. 2015. Going through, going around: a study on individual avoidance of groups. *Visualization and Computer Graphics, IEEE Transactions on PP*, 99, 1–1.
- FIORINI, P., AND SHILLERT, Z. 1998. Motion planning in dynamic environments using velocity obstacles. *International Journal of Robotics Research* 17, 760–772.
- GERAERTS, R. 2010. Planning short paths with clearance using explicit corridors. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 1997–2004.
- GOLAS, A., NARAIN, R., CURTIS, S., AND LIN, M. 2014. Hybrid long-range collision avoidance for crowd simulation. *Visualization and Computer Graphics, IEEE Transactions on* 20, 7 (July), 1022–1034.
- GUY, S. J., CHHUGANI, J., CURTIS, S., DUBEY, P., LIN, M., AND MANOCHA, D. 2010. Pedestrians: A least-effort approach to crowd simulation. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SCA '10, 119–128.
- HELBING, D., FARKAS, I., AND VICSEK, T. 2000. Simulating dynamical features of escape panic. *Nature* 407, 6803, 487–490.
- KALLMANN, M. 2010. Shortest paths with arbitrary clearance from navigation meshes. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SCA '10, 159–168.
- KAPADIA, M., SINGH, S., HEWLETT, W., AND FALOUTSOS, P. 2009. Egocentric affordance fields in pedestrian steering. In *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games*, ACM, New York, NY, USA, I3D '09, 215–223.
- KAPADIA, M., BEACCO, A., GARCIA, F., REDDY, V., PELECHANO, N., AND BADLER, N. I. 2013. Multi-domain real-time planning in dynamic environments. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ACM, New York, NY, USA, SCA '13, 115–124.
- KARAMOUZAS, I., GERAERTS, R., AND OVERMARS, M. 2009. Indicative routes for path planning and crowd simulation. In *Proceedings of the 4th International Conference on Foundations of Digital Games*, ACM, 113–120.
- LAMARCHE, F., AND DONIKIAN, S. 2004. Crowd of virtual humans: a new approach for real time navigation in complex and structured environments. *Computer Graphics Forum* 23, 3, 509–518.
- ONDŘEJ, J., PETTRÉ, J., OLIVIER, A.-H., AND DONIKIAN, S. 2010. A synthetic-vision based steering approach for crowd simulation. *ACM Trans. Graph.* 29, 4, 123:1–123:9.
- PETERS, C., ENNIS, C., AND C., O. 2009. Modeling groups of plausible virtual pedestrians. *Computer Graphics and Applications, IEEE* 29, 4, 54–63.
- PETTRÉ, J., CIECHOMSKI, P. D. H., MAÏM, J., YERSIN, B., LAUMOND, J.-P., AND THALMANN, D. 2006. Real-time navigating crowds: scalable simulation and rendering. *Computer Animation and Virtual Worlds* 17, 3-4, 445–455.
- REYNOLDS, C. W. 1987. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Comput. Graph.* 21, 4 (Aug.), 25–34.
- SCHADSCHNEIDER, A. 2001. Cellular automaton approach to pedestrian dynamics-theory. In *Pedestrian and Evacuation Dynamics*, 75–86.
- SHAO, W., AND TERZOPOULOS, D. 2005. Autonomous pedestrians. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, ACM, 19–28.
- SUD, A., GAYLE, R., ANDERSEN, E., GUY, S., LIN, M., AND MANOCHA, D. 2007. Real-time navigation of independent agents using adaptive roadmaps. In *Proceedings of the 2007 ACM Symposium on Virtual Reality Software and Technology*, ACM, New York, NY, USA, VRST '07, 99–106.
- VAN DEN BERG, J., LIN, M., AND MANOCHA, D. 2008. Reciprocal velocity obstacles for real-time multi-agent navigation. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 1928–1935.
- WHITTLE, M. W. 2003. *Gait analysis: an introduction*. Elsevier.
- ZIPF, G. K. 1949. *Human Behaviour and the Principle of Least-Effort*. Addison-Wesley, Cambridge MA.