


Adding Sociality to Virtual Pedestrian Groups

Norman Jaklin
Utrecht University*

Angelos Kremyzas
Utrecht University*

Roland Geraerts
Utrecht University*



Abstract

We present a new method, *Social Groups and Navigation* (SGN), to simulate the walking behavior of small pedestrian groups in virtual environments. SGN is the first method to simulate group behavior on both global and local levels. We define quantitative metrics to measure the *coherence* and the *sociality* of a group based on existing empirical data of real crowds. SGN is designed to let groups stay in coherent and social formations without explicitly modeling such formations. For groups of four, SGN generates between 13% and 53% more socially-friendly behavior than existing methods, measured over the lifetime of a group in the simulation. For groups of three, the gain is between 15% and 31%, and for groups of two, the gain is between 1% and 4%. SGN can be used with any existing global path planner and any existing path following method. Experiments show that SGN enables the simulation of thousands of agents in real time. Due to its flexible design, it can be easily integrated into a larger crowd simulation framework.

CR Categories: I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Intelligent agents I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: virtual humans, social-group behavior, autonomous agent navigation

1 Introduction

The simulation of pedestrian crowds has become increasingly important in safety training for mass events and evacuation scenarios, urban city planning, cgi-enhanced movies and video games.

* {n.s.jaklin, a.kremyzas, r.j.geraerts}@uu.nl

This research was funded by the COMMIT/ project: www.commit-nl.nl

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.

VRST '15, November 13 – 15, 2015, Beijing, China.
© 2015 ACM. ISBN 978-1-4503-3990-2/15/11...\$15.00
DOI: <http://dx.doi.org/10.1145/2821592.2821597>

Over the past decades, a wide range of models have been developed that simulate a crowd as individual agents or as one entity based on flow dynamics. Less attention has been paid to the simulation of small social groups, although empirical research shows that a high percentage of crowd members walk in groups in urban environments and public places [James 1953; Coleman and James 1961; Moussaïd et al. 2010].

Existing methods model explicit formations to keep groups coherent [Kimmel et al. 2012] and socially-friendly [Karamouzas and Overmars 2012; Wu et al. 2013]. Such formations have been observed in real crowds [Moussaïd et al. 2010], but they are not strictly kept at all times due to the wide range of factors that influence a group's walking behavior. We therefore believe that explicitly modeling such formations may yield artificial-looking behavior. Groups may lack flexibility and put too much emphasis on maintaining an explicit formation. For instance, groups might not be able to temporarily split and instead take unrealistic detours.

Contributions. We present a novel method named *Social Groups and Navigation* (SGN) to simulate the walking behavior of small pedestrian groups. SGN is based on the social-force model by [Moussaïd et al. 2010] and the vision-based collision-avoidance method by [Moussaïd et al. 2011], which we have modified and extended to yield more coherent and socially-friendly behavior. We do not explicitly model social formations. We instead introduce quantitative metrics to measure the *coherence* and *sociality* of small groups. We use these metrics to let formations emerge from the group members' attempts to stay coherent and social. The generated behavior is more flexible and diverse than with existing methods. SGN allows groups to temporarily split to avoid dynamic obstacles such as other agents, and groups automatically re-organize themselves when coherence is lost. Thus, SGN handles social behavior on both global and local levels of a crowd simulation framework. Furthermore, SGN can be used to simulate thousands of agents in real time. Our method does not depend on a particular representation of the environment. Any representation that allows real-time path planning with clearance from obstacles is a feasible choice, e.g. [Geraerts 2010; Kallmann 2010; Oliva and Pelechano 2015]. SGN can be easily integrated into existing crowd simulation frameworks that independently support global planning, route following and local behavior such as collision avoidance, e.g. [Curtis et al. 2014; van Toll et al. 2015].

2 Related Work

For a general overview of crowd simulation, we refer the reader to the books by [Thalmann and Musse 2013] and [Pelechano et al. 2008]. In this section, we focus on selected work related to our SGN method.

Early work on social groups did not involve coherence or social formations. [Musse and Thalmann 1997] described a model that uses a small set of parameters such as *interests*, *emotional status*, and *domination* for the agents. The focus is on social relationships between groups and their members.

[Qiu and Hu 2010] presented a model to simulate pedestrian groups based on utility theory and social-comparison theory. Agents are allowed to switch between groups, but coherence and socially-friendly formations are not modeled.

Other models focus on coherence. [Kamphuis and Overmars 2004] presented a method to simulate large coherent groups such as military armies. It handles both global planning and local steering. Socially-friendly formations are not supported, and coherence is not re-established when it is lost. [Kimmel et al. 2012] presented an extension of the *Velocity Obstacle* (VO) [Fiorini and Shiller 1998] approach to simulate coherent groups. A geometrical *Loss of Communication Obstacle* (LOCO) is used, which can be combined with a VO to generate collision-free movement for small groups. Coherence is kept by preventing agents to move farther away from their groups than a threshold distance. There is no explicit formulation of socially-friendly formations, and the method works only locally as an extension of the VO method and its reciprocal variants, e.g. [van den Berg et al. 2009]. [Park et al. 2012] presented a model that considers higher-level social interactions between the group members. It assigns a *leader* to each group, and it handles group-coordination strategies based on common ground theory. [Huang et al. 2014] present a path-planning method to simulate coherent and persistent groups. It is based on the *Local Clearance Triangulation* by [Kallmann 2010], and it handles groups as deformable shapes. Deformations as well as splitting and merging actions of a group influence the path costs.

Other models simulate socially-friendly formations. [Moussaïd et al. 2010] use video recordings of urban areas to collect empirical data of pedestrian crowds. They also describe a social-force model to simulate small pedestrian groups. Our SGN method is based on this social-force model. Inspired by [Moussaïd et al. 2010], [Karamouzas and Overmars 2012] presented a velocity-based approach. The method explicitly models socially-friendly formations, and it optimizes a cost function to maintain group coherence and guarantee collision-free movement. [Wu et al. 2013] combine the work by [Karamouzas and Overmars 2012] with the vision-based steering approach by [Ondřej et al. 2010]. They compare the distortion, dispersion, and out-of-formation metrics of their simulation with data from a real crowd.

3 Preliminaries

3.1 Basic Settings

We assume that SGN is used in the context of a larger crowd simulation framework. We have implemented it within the framework described by [van Toll et al. 2015]; see Figure 1.

Assume we are given k groups of agents \mathcal{G}_i , $1 \leq i \leq k$. We assume group sizes $|\mathcal{G}_i|$ of 2 through 4, which corresponds to observations made in real crowds [Moussaïd et al. 2010]. Note that these sizes are not a hard constraint, and SGN can be easily modified to simulate bigger groups; see Section 4.6.3. The method is designed in a modular way, and it can be used to also simulate individuals by switching off the corresponding group-related parts

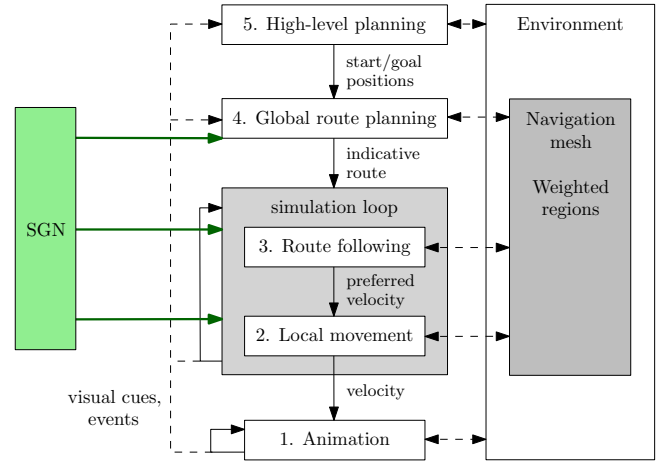


Figure 1: Example of a multi-level crowd-simulation framework [van Toll et al. 2015] into which our SGN method can be integrated.

of the method. It also allows to simulate mixed scenarios with groups and individuals. For ease of explanation, we assume that all groups are present at the start of the simulation. However, our method can be easily modified to let groups enter the simulation at a later point in time.

For $1 \leq i \leq k$ and $1 \leq j \leq |\mathcal{G}_i|$, we denote by A_{ij} the j th member of group \mathcal{G}_i , which is represented as a disc with radius r_{ij} (in m) and a mass $m_{ij} = 320r_{ij}$ (in kg), following the definition by [Moussaïd et al. 2011]. By x_{ij} , we denote the center point of the disc that represents A_{ij} , and we refer to it as the agent's position. Each agent has a *personal space radius* $r'_{ij} \geq r_{ij}$ and a *preferred speed* s_{ij} . Each group has a *preferred group speed* $s_i = \min_{1 \leq j \leq |\mathcal{G}_i|} s_{ij}$, which is the smallest preferred speed of its members.

Each agent A_{ij} has a *field of view* (FOV), which is a circular segment centered in x_{ij} with a viewing distance d_{ij} and a viewing angle of Φ_{ij} . We say that an agent $A_{i'j'}$ is *visible* to an agent A_{ij} , if the FOV of A_{ij} contains at least one point of

the disc that represents $A_{i'j'}$. We assume that $d_{ij} \geq 2 \sum_{j'=1}^{|\mathcal{G}_i|} r'_{ij'}$

to ensure that an agent can see all its group members when they are lined up in front of the agent, with the personal spaces of any two consecutive agents overlapping in at least one point (Figure 2). This is important for (re-)establishing group coherence; see Section 4.3.

Each group \mathcal{G}_i has a goal area G_i . We assume that \mathcal{G}_i has a feasible global route to G_i . A feasible route can be computed with any existing path planning method that ensures clearance from static obstacles, e.g. [Geraerts 2010; Kallmann 2010; Oliva and Pelechano 2015]. To ensure collision-free movement for *all* group members, the global route should keep clearance from obstacles that corresponds to the *largest* disc radius of all group members.

3.2 Overview of the SGN Method

A group that enters the simulation first establishes its coherence by letting all members walk towards the group leader. All coherent groups walk towards their goal along a shared global path. While walking, social forces try to make the group members stay coherent and social. Whenever coherence is lost, a group re-establishes its coherence by letting the leader wait for its fellow members as soon as the local crowd density around the leader is low.

We now give an overview of the initialization step in Section 3.2.1,




Figure 2: Example of a group of 3 agents that are lined up. Agent A_{i1} can perceive its fellow group members because its viewing distance d_{i1} is larger than the sum of the personal space diameters of all group members.

and we continue with the simulation loop in Section 3.2.2.

3.2.1 SGN Initialization

For each group \mathcal{G}_i , we perform the following actions:

- Set \mathcal{G}_i to *coordination mode* (Section 4.3).
- Let an arbitrary member of \mathcal{G}_i be the *leader* L_i (Section 4.1).
- Compute a route π_{ij} to L_i for each member $A_{ij} \neq L_i$.
- Compute a route π_i from L_i to the goal area G_i .

3.2.2 The SGN Simulation Loop:

For each group \mathcal{G}_i in *coordination mode* (Section 4.3), we perform the following actions:

- Compute a preferred velocity for each $A_{ij} \neq L_i$ to move along π_{ij} with speed s_{ij} . Any existing path-following method can be used.
- Pass the preferred velocities to a modified version of the collision-avoidance method by [Moussaïd et al. 2011] (Section 4.5).
- Check for each *waiting* member whether there is a fellow agent in its personal space. If so, set that agent to a *waiting* state, too.
- If all group members are in a *waiting* state, then set \mathcal{G}_i to *group-walking mode* (Section 4.4).

For each group \mathcal{G}_i in *group-walking mode* (Section 4.4), we perform the following actions:

- (Re-)assign the role of L_i to the group member that is closest to G_i , measured via the curve-length distance along π_i (Section 4.1).
- Determine the current last member l_i of the group, which is farthest away from G_i , measured via the curve-length distance along π_i (Section 4.1).
- Compute a preferred velocity for each member along the group's global path π_i with speed s_i . Any existing path-following method can be used here.
- Pass the preferred velocities to a modified version of the collision-avoidance method by [Moussaïd et al. 2011] (Section 4.5).

- Compute the acceleration for each agent using a modified version of the social-force model by [Moussaïd et al. 2010] (Section 4.4).
- If \mathcal{G}_i is not *coherent* (Section 4.2) and the density around L_i is smaller than 0.7 pedestrians per m^2 [Fruin 1971], then set \mathcal{G}_i to *coordination mode* (Section 4.3).

4 The Social Groups and Navigation Method

4.1 Leader and Last Member

For each group \mathcal{G}_i , we define a *leader* L_i and a *last member* l_i . These roles are re-assigned at the end of each simulation cycle. L_i is defined to be the group member that is closest to G_i , measured via the curve-length distance along π_i . The last member is defined as the member that is farthest away from G_i ; see Figure 3 for an example. The only exception is in the initialization phase. Here, no global path has been computed yet, and the role of the leader is assigned to an arbitrary member.

Note that the global path serves as an *indicative route* [Karamouz et al. 2009], and the agents are in general *not* located exactly on that route. It depends on the path-following method what points on the route are used to determine the leader and the last member. One option is to define a *reference point* on the global path for each agent, e.g. the point on the global path that is closest to an agent's position [Jaklin et al. 2013].

4.2 Coherence and Sociality

We define the *coherence* and *sociality* of G_i in the following way:

Definition 1. Let \mathcal{G}_i be a group with leader A_{ij} and last member $A_{ij'}$. We say that \mathcal{G}_i is *coherent* iff $\|x_{ij} - x_{ij'}\| \leq d_{ij'} + r_{ij}$.

In other words, a group is coherent when at least one point of the disc that models the leader can potentially be seen by the last member. Note that this does not reflect whether the leader is actually *inside* the FOV of the last member. As long as their distance is within the defined range, the group is coherent, even when the last member is not looking in the leader's direction.

We define a *social threshold distance* d_{social} . It is a maximum distance that two group members are allowed to keep from each other while still being able to socially interact. This distance is based on empirical observations [Zipf 1949; Costa 2010; Moussaïd

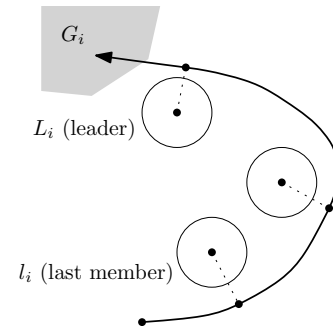


Figure 3: Example of a group \mathcal{G}_i with a leader L_i and a last member l_i , following an indicative route to a goal area G_i . The discs resemble the group members themselves, and we do not display their personal space discs. The dotted line segments indicate the distance from each agent to its reference point on the indicative route.

et al. 2010; Fridman et al. 2013]. It should not be larger than the minimum of the viewing distances of all agents, i.e. $\forall i \in [1, \dots, k]$ $\forall j \in [1, \dots, |\mathcal{G}_i|]$ $d_{social} \leq d_{ij}$.

Definition 2. We say that a group \mathcal{G}_i is in a partially social configuration iff $\forall j \in [1, \dots, |\mathcal{G}_i|] \exists j' \in [1, \dots, |\mathcal{G}_i|]$, $j \neq j'$, such that A_{ij} and $A_{ij'}$ are mutually visible and $\|x_{ij} - x_{ij'}\| \leq d_{social} + r_{ij} + r_{ij'}$.

Definition 3. We say that a group \mathcal{G}_i is in a totally social configuration iff \mathcal{G}_i is partially social and $\forall j \in [1, \dots, |\mathcal{G}_i|] \forall j' \in [1, \dots, |\mathcal{G}_i|]$, $j \neq j'$, A_{ij} and $A_{ij'}$ are mutually visible.

A group is partially social when each member has at least one mutually visible other member within the social threshold distance, and it is totally social when, in addition, all members are mutually visible. These definitions are based on criteria that were defined in previous studies on pedestrian groups [James 1953]. In these studies, clear social interactions among group members (talking, laughter, smiles, gesticulation) are used to identify people belonging to the same group [Moussaïd et al. 2010].

4.3 Coordination mode

When a group loses its coherence, it enters *coordination mode*. In this mode, the members will gather around L_i to (re-)establish coherence. L_i enters a *waiting* state and does not move until coherence is established. *Coordination mode* is also the default mode of each new group in the simulation. Each $A_{ij} \neq L_i$ in *coordination mode* computes a route to L_i and starts following it. Any path-planning method that guarantees clearance from obstacles is sufficient [Geraerts 2010; Kallmann 2010; Oliva and Pelechano 2015]. In the same way, any route-following method can be used that computes a preferred velocity for each agent. In *coordination mode*, this preferred velocity is then passed to a collision-avoidance method; see Section 4.5.

Any $A_{ij} \neq L_i$ follows its route to the L_i until it detects a fellow member that is in a *waiting* state. When *coordination mode* starts, only L_i is in a *waiting* state. Members that are close to L_i enter a *waiting* state, too. We use an agent's personal space to determine whether another agent is sufficiently close: At the end of each simulation cycle, each *waiting* member of a group in *coordination mode* checks whether there is a *non-waiting* member in its personal space. If so, that *non-waiting* member switches to a *waiting* state, too. Since we assume the viewing distance d_{ij} of each agent A_{ij} to

be at least $2 \sum_{j'=1}^{|\mathcal{G}_i|} r_{ij'}$ (see Section 3), the group will always be coherent as soon as all members have switched to a *waiting* state. This ensures that we can safely set the group to *group-walking mode* when there are no *non-waiting* members left at the end of a simulation cycle.

4.4 Group-walking mode

In this mode, each group \mathcal{G}_i moves along π_i to G_i . Any route-following method can be used that takes a guidance path as an input and computes a preferred velocity for each agent. This preferred velocity is then passed to a collision-avoidance method; see Section 4.5. Afterwards, the preferred velocity is passed to a social-force model that maintains group-coherence and sociality; see Section 4.6.

After the group has moved, we check whether it is coherent according to Definition 1. If not, the group needs to re-establish its coherence. In real-life, we expect the leader to wait for its fellow group members in a non-congested area of the environment. Thus, in order to prevent a leader from stopping in the middle of a highly

dense situation, we check whether the local crowd density around the leader is smaller than a threshold value of 0.7 agents per m^2 . This value is based on the *Pedestrian Level Of Service* (PLOS) system proposed by [Fruin 1971]. Only when the group lost its coherence and when the local crowd density around the leader is small, we set the group back to *coordination mode*.

4.5 Collision-Avoidance

We use the vision-based collision-avoidance method by [Moussaïd et al. 2011] with some modifications. We keep the following core concepts as proposed in the original method:

Let v_{ij} be the preferred velocity of agent A_{ij} as computed by the path-following algorithm that is used. Let α_0 be the corresponding angle of v_{ij} measured against the agent's line of sight. Let O_{ij} be the last visible point in A_{ij} 's FOV that lies in the direction α_0 . Let $\alpha \in [-\frac{\Phi_{ij}}{2}, +\frac{\Phi_{ij}}{2}]$ be a candidate angle direction, and let Ω_α be the last visible point in A_{ij} 's FOV that lies in the direction of α . Let T_α be the point in the direction of α that is the last collision-free position within the agent's FOV. Figure 4 shows an example of the situation. The desired direction is then defined as

$$\alpha_{des} = \underset{\alpha \in [-\frac{\Phi_{ij}}{2}, +\frac{\Phi_{ij}}{2}]}{\operatorname{argmin}} d(\alpha),$$

where $d(\alpha) = \sqrt{d_{ij}^2 + f(\alpha)^2 - 2d_{ij}f(\alpha)\cos(\alpha_0 - \alpha)}$.

In the original method, the term $f(\alpha)$ is defined as the distance from the agent's position to T_α . If no collision occurs within the distance of d_{ij} , then T_α coincides with Ω_α , and $f(\alpha)$ is therefore set to d_{ij} . For our method, we modify the definition of $f(\alpha)$: We let F_α be the perpendicular foot of O_{ij} on the straight-line segment between x_{ij} and Ω_α (Figure 4). We then define

$$f(\alpha) = \min(|x_{ij}T_\alpha|, |x_{ij}F_\alpha|).$$

The moment when the next directional change occurs should not solely be based on the impending collisions, but also on the distance to O_{ij} . In other words, an agent should also change its direction when it reaches a point where moving on in its current direction would increase the distance to O_{ij} , even with no impending collisions. The point where this happens is F_α . Without that option, an agent might 'overshoot' in the desired direction.

Now that α_{des} is computed, we can compute the desired speed s_{des} . From the original method, we adopt the concept of a relaxation time τ . This relaxation time ensures that an agent chooses its speed such that there is enough time to avoid a collision within the given time frame. Let s_{ij} be the agent's preferred speed, and let d_{col} be the distance between the agent and the first collision in the direction of




Figure 4: Example of the situation during the modified collision-avoidance method for a candidate angle α . Here, the point F_α is closer to x_{ij} than T_α . Thus, $f(\alpha)$ is set to $|x_{ij}F_\alpha|$.

α_{des} . We define

$$s_{des} = \min(s_{ij}, d_{col}/\tau).$$

This ensures that the agent moves at its preferred speed when there are no impending collisions in the given time frame, and it slows down accordingly when needed. Finally, let v_{des} be the resulting desired velocity in the direction of α_{des} and scaled to the desired speed s_{des} .

4.6 Social-Force Model

We apply social forces to each group in *group-walking mode* after the collision-avoidance step. The social forces are based on the model by [Moussaïd et al. 2010], with some modifications. We will first explain the model and then discuss what details have been modified compared to the original method.

Given an arbitrary agent A_{ij} , its desired velocity v_{des} , and the actual velocity v from the previous simulation step, we compute the acceleration $\frac{dv}{dt}$ in the following way:

$$\frac{dv}{dt} = \frac{v_{des} - v}{\tau} + \frac{1}{m} \sum_{u=1}^k \sum_{v=1}^{|\mathcal{G}_u|} f_{uv} + \frac{1}{m} \sum_{w=1}^{\mathcal{W}} f_w + \frac{f_{group}}{m}, \quad (1)$$

where f_{uv} is a repelling force to avoid another agent A_{uv} (Section 4.6.1), f_w is a repelling force to avoid one of the \mathcal{W} wall segments in the environment (Section 4.6.2), and f_{group} is a group force to maintain coherent and socially-friendly formations (Section 4.6.3).

4.6.1 Physical-contact Force with another Agent

The force f_{uv} is applied to agent A_{ij} when there is contact with agent A_{uv} . By $dist(A_{ij}, A_{uv})$, we denote the Euclidean distance between A_{ij} and A_{uv} . Let $n(A_{ij}, A_{uv}) = \frac{x_{ij} - x_{uv}}{dist(A_{ij}, A_{uv})}$ be the unit vector pointing from A_{uv} to A_{ij} . Let S be a global parameter that defines the strength of the force [Moussaïd et al. 2011]. We then define the force as follows:

$$f_{uv} = S \cdot \max\left(0, r_{ij} + r_{uv} - dist(A_{ij}, A_{uv})\right) \cdot n(A_{ij}, A_{uv}).$$

4.6.2 Physical-contact Force with Obstacles

The force f_w is applied to agent A_{ij} when there is contact with a wall segment w in the environment. By $dist(A_{ij}, w)$, we denote the Euclidean distance between A_{ij} and w . Let $n(A_{ij}, w)$ be a unit-length vector that is perpendicular to w and points from w to A_{ij} . Let S be the global force-strength parameter as described in Section 4.6.1. We then define the force f_w as follows:

$$f_w = S \cdot \max\left(0, r_{ij} - dist(A_{ij}, w)\right) \cdot n(A_{ij}, w).$$

4.6.3 Group Force

The group force f_{group} is defined as $f_{group} = f_{vis} + f_{att}$, where f_{vis} is a deceleration force that represents the desire of A_{ij} to keep its fellow group members in its FOV, and f_{att} is an attractive force that represents the desire of A_{ij} to maintain group coherence.

To define f_{vis} , let $\theta_{ij'}, j' \neq j$, be the minimum rotation angle (in degrees) that is required to let the position $x_{ij'}$ of agent $A_{ij'}$ be inside agent A_{ij} 's FOV. Let $\theta = \max_{1 \leq j' \leq |\mathcal{G}_i|} \theta_{ij'}$ be the maximum of the minimum rotation angles. Furthermore, let S_{vis} be a global

parameter that defines the strength of f_{vis} [Moussaïd et al. 2010]. We then define f_{vis} as follows:

$$f_{vis} = -S_{vis} \cdot \theta \cdot v_{des}$$

This means that we scale the desired velocity v_{des} by the rotation angle θ and the strength parameter S_{vis} in negative direction of v_{des} to compute the first part of the group force.

The force f_{att} describes agent A_{ij} being attracted to the centroid

$$C_i = \frac{1}{|\mathcal{G}_i|} \sum_{1 \leq j \leq |\mathcal{G}_i|} x_{ij}$$

of the group \mathcal{G}_i (viewed as a given set of points) to maintain group coherence. Let $dist(A_{ij}, C_i)$ be the distance from the agent to the centroid. Similar to [Moussaïd et al. 2010], we define a threshold distance $d = 0.5 \cdot (|\mathcal{G}_i| - 1)$, such that A_{ij} is attracted to C_i as soon as its distance to C_i exceeds d . Let $n(A_{ij}, C_i)$ be the vector pointing from C_i to A_{ij} , normalized to unit length. Furthermore, let S_{att} be a global parameter that defines the strength of f_{att} [Moussaïd et al. 2010]. We then define f_{att} as follows:

$$f_{att} = \begin{cases} S_{att} \cdot n(A_{ij}, C_i), & \text{if } dist(A_{ij}, C_i) \geq d \text{ and } v_{des} \neq \mathbf{0} \\ \mathbf{0}, & \text{otherwise} \end{cases}$$

In the above definition, we check whether the desired velocity v_{des} given by the collision-avoidance method is $\mathbf{0}$. If so, this means that the agent has reached its goal, and we therefore let the attraction force be $\mathbf{0}$, too. This yields an overall group force of $\mathbf{0}$, and it disables social behavior for agents that have reached their goal.

According to [Costa 2010], big social groups in real-life tend to split up into smaller sub-groups of up to 3 members. Our SGN method can be easily adjusted to account for this behavior in the computation of the visual group force f_{vis} : We can split up each group into subgroups of at most 3 members. Instead of computing f_{vis} with respect to all group members, only the members of agent A_{ij} 's sub-group are taken into consideration. All other steps of the method remain unchanged.

4.6.4 Differences to the Original Model

In the original social-force model by [Moussaïd et al. 2010], the acceleration term for agent A_i is defined as

$$\frac{dv_i}{dt} = f_i^0 + f_i^{wall} + \sum_j f_{ij} + f_i^{group}.$$

Here, f_i^0 is an attractive force to move agent A_i in a particular direction at a preferred speed, f_i^{wall} is a repulsive force to avoid static obstacles, f_{ij} is a repelling force to avoid physical contact with another agent A_j from a different group. The resulting behavior is reactive and lacks anticipation. To add a more predictive avoidance behavior within our SGN method, we have replaced the above forces by the avoidance forces of [Moussaïd et al. 2011]; see Equation 1 in Section 4.6.

Another modification is that we use the centroid C_i of the group \mathcal{G}_i when computing the group force f_{group} . In the original method, the center of mass of the group is used instead of the centroid. We assume that a variation in mass among the group members should not have an effect on the group force, which is why we consider the centroid a better choice.

Similarly, we modified the computation of the force f_{vis} : In the original method, the force is defined via the required rotation angles for each agent to keep the center of mass of the group in its

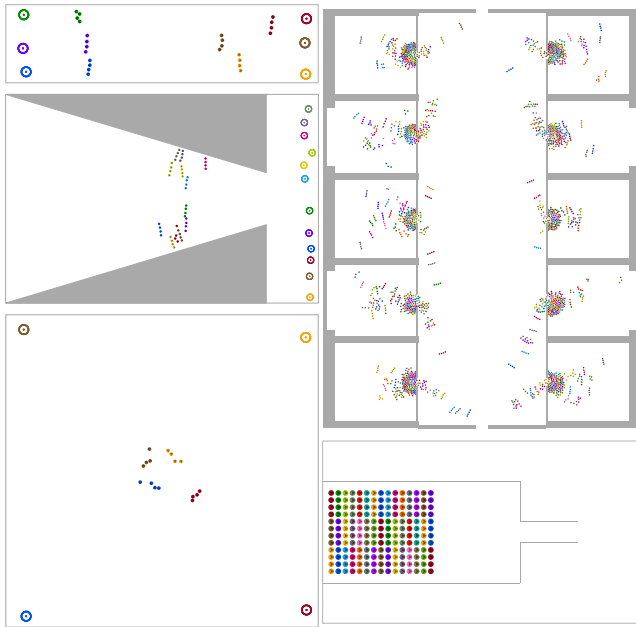


Figure 5: The scenarios we used for our experiments, shown with groups of 4. Left side from top to bottom: bidirectional corridor, bottleneck, and corners. Right side from top to bottom: building evacuation, and room evacuation. Small discs indicate the agents, grouped by color, and large discs indicate the corresponding goal areas.

FOV. Instead, we define the force via the required rotation angles to keep the group members themselves in an agent’s FOV. Again, we believe that a variation in mass should not have an effect on this step. Furthermore, the original method does not guarantee that group members effectively slow down when a fellow member is left behind in dense situations, which our modification does.

Another change in the force f_{vis} is that we use the desired velocity v_{des} that already takes predictive avoidance behavior into account. In the original method, the actual velocity of an agent is used here, which lacks anticipation.

Finally, we changed the repulsive forces between agents: In the original method, the group force f_i^{group} contains a repulsive term to model the interaction between members of the same group. In our SGN model, we skip this term. Physical contact between agents are generally resolved by our definition of f_{uv} in Equation 1, independent of whether the agents are from the same group or not.

5 Experiments

5.1 Experimental Setup

We have tested and validated our SGN method on a PC with an Intel Core i7 860 processor with 2.8 GHz, an Nvidia GeForce GTX 285 video card and 8 GB of RAM, running Windows 7 Ultimate 64bit. We have used one single core for all experiments.

We integrated the method into the crowd simulation framework described by [van Toll et al. 2015]. For each agent, we used a radius of 0.24 m, an FOV of $\Phi_{ij} = \pi$ with maximum viewing distance of 10 m. The personal space radius and the social threshold distance were set to 1 m each. Following [Weidmann 1992], we used a normal distribution with a mean of 1.34 m/s and a standard deviation of 0.26 m/s to randomly choose the preferred speed for each agent. Each goal area was modeled as a disc with radius 0.6 m.

The relaxation time τ used in the social-force model was set to 0.5 s. Following [Moussaïd et al. 2011], the strength parameter S of the physical forces was set to 5000. Following [Moussaïd et al. 2010], the parameter S_{att} was set to 3. The parameter S_{vis} , which was set to 4 in [Moussaïd et al. 2010] yielded undesired stop-and-go behavior among group members, and we instead determined a value of 1 via preliminary experiments based on visual inspection. Furthermore, we set the time for one simulation step to 0.1 s.

We tested our method with group sizes of 2, 3, and 4 in five different scenarios: *bidirectional corridor*, *bottleneck*, *corners*, *building evacuation*, and *room evacuation*. The scenarios are displayed in Figure 5.

Bidirectional corridor features a 20 m long corridor that is 10 m wide. Three groups are placed on each end of the corridor, and each group has its goal areas at the opposite end of the corridor. We use this scene to test whether groups stay coherent and in socially-friendly formations when they encounter other groups moving in the opposite direction.

Bottleneck features a 50 m long corridor that linearly decreases in width towards the right side. On the left, the corridor is 40 m wide, and on the right it is 10 m wide. Twelve groups are placed on the left end and have their goal areas on the right end. We use this scene to test whether groups stay coherent and in socially-friendly formations when the environment becomes more narrow and crowd density increases.

Corners features an empty square room with four social groups. Each group is placed near a different corner and has its goal position near the opposite corner of the room. We use this scene to test whether groups stay coherent and in socially-friendly formations when having to cross the center point of a room with other groups approaching from different directions.

Building evacuation features a building that spans an area of $95 \text{ m} \times 128 \text{ m}$. The building has ten rooms that are connected via one large corridor. The corridor has an exit at each end. A total of 490 groups is placed in the rooms, and each group has to leave the building through the nearest exit. The members of a group are all placed at random positions in the same room. We use this scene to test whether groups stay coherent and in socially-friendly formations in a high-density evacuation situation.

Room evacuation features a room with one exit, and a crowd of 180 agents subdivided into groups of varying size. The agents have to evacuate the room through the exit. This experimental setup was proposed by [Köster et al. 2014]. It is based on a controlled laboratory experiment performed by [Liddle et al. 2011]. We use this scene to test whether our SGN method generates group behavior that is in line with empirical data, and what effect the group size has on evacuation times.

5.2 Effects of SGN on Coherence and Sociality

In a first set of experiments, we compared our SGN method against the methods by Moussaïd et al. [2010; 2011]. To this end, we integrated both the collision-avoidance method [Moussaïd et al. 2011] and the social-force model [Moussaïd et al. 2010] into the *local movement* layer of the framework (Figure 1) by [van Toll et al. 2015].

The goal was to test whether SGN with its additions to the combined work by Moussaïd et al. yields group behavior that better reflects real-life situations than the original methods. An overall assumption is that agents do not switch groups during the simulation. Thus, real-life behavior in a corresponding situation means that each person tries to stay in coherent and socially-friendly formations as much as possible while approaching the goal area. We therefore compared the frequency of coherence and sociality in our simulated groups for the SGN method and for the work by Moussaïd et al. We measured the percentage of simulation steps

over the *lifetime* of a group in which it is coherent according to Definition 1. By *lifetime*, we refer to the number of simulation steps that it takes a group to reach its goal area. Similarly, we measured the percentage of simulation steps over the lifetime of a group in which it is in a partially-social and totally-social formation according to Definitions 2 and 3, respectively. We ran each scenario 100 times and took the average coherence and sociality over all runs.

Table 6 (top) shows the average coherence (%) for both methods and varying group sizes in the first four scenarios. Figures 6 (center) and 6 (bottom) show the average partial and total sociality (%), respectively.

The results show that our SGN method improves over the work by Moussaïd et al. in all cases with respect to partial and total sociality. Regarding coherence, our method improves in all cases except the *bottleneck* scenario with groups of 3. In that scenario, coherence is lost in one single run for our SGN method. A Welch's t-test on the difference between the two coherence results (for SGN and Moussaïd et al.) yielded a p -value of 0.3198, and the difference is thus not considered statistically significant.

5.3 Evacuation Times

In a second set of experiments, we used the *room evacuation* scenario (Figure 5, bottom-right) with 180 agents. We measured the evacuation times achieved by our SGN method, by Moussaïd et al. [2010; 2011] and by [Köster et al. 2014]. We set the radius of each agent to 0.2 m. All other settings were kept as described in Section 5.1. Since the constrained space for this scenario does not allow for much variation in the initial spacial distribution of the groups, we used a fixed initial configuration for the 180 agents. With no randomness left, we ran our SGN method once per group size and measured the total time needed to evacuate the room.

Figure 7 shows the result of this experiment. The corresponding real-life experiment by [Liddle et al. 2011] was performed with 180 individuals, for which a total evacuation time of 80 s was reported. There is no corresponding ground truth data for bigger group sizes. However, according to empirical data obtained by [Xu and Duh 2010], the evacuation times should increase when the group size increases. Indeed, both SGN and the method by Köster et al. show this trend. With the methods by Moussaïd et al., a decrease in evacuation times can be observed for groups of 4, which contradicts the empirical observations.

In addition to the group size, we tested the effect that the radius of an agent's disc has on evacuation time when using SGN. To this end, we repeated the scenario four times with all radii increased, ranging from 0.21 m up to 0.24 m with a step size of 0.01 m. Furthermore, we ran a variant of this scenario with mixed radii that were randomly chosen in the range of 0.20 m to 0.24 m for each agent. Figure 8 shows the result of these experiments. We conclude that an increase in the radius increases the evacuation times for all group sizes. This is an expected result because higher radii yield less free space in the environment, which increases the overall crowd density. For groups of 2 and radii of 0.23 and 0.24, we observed high evacuation times due to an increased crowd congestion in these runs. Increased crowd congestion could also be observed for single runs with other group sizes and even when using only individual agents. This effect occurs due to the highly competitive nature of the collision avoidance method, which does not guarantee coordination of crowds in high-density scenarios but rather makes agents try to exploit gaps in the crowd whenever possible.

5.4 Performance

We tested the performance of our SGN method in an extended variant of the *room evacuation* scenario, which consists of eleven copies of the scenario, i.e. eleven rooms as displayed in Figure 5,

bottom-right. Each room is initially occupied by 180 agents, yielding a total of 1980 agents in this stress-test scenario. The agents are subdivided into groups, and each group has to evacuate the room it is starting in. We ran the scenario 100 times for group sizes of 1, 2, 3, 4, and mixed sizes, and we measured the average time needed to compute one simulation step.

Figure 9 shows the average time per simulation step (top) and frame rates (bottom) we achieved for a *serial* and *parallel* execution our method. For the parallel execution, we used 4 CPU cores and a total of 8 threads. The results show that the average running times are all close to each other for the varying group sizes, and mixed group sizes yield intermediate running times. For all group sizes, our SGN method only yields a small increase in average running times over the simulation of individual agents. When executing the method in parallel, one simulation step is performed about 4.5 times as fast as with a serial execution. For all tested group sizes, we achieved an average rate of slightly less than 20 steps per second. Since we set the time for one simulation step to 0.1 s, we can conclude that our SGN method achieves real-time performance for large numbers of agents when using parallel computation.

6 Limitations

While our SGN method generates coherent and socially-friendly group behavior for a large number of agents in real time, it has some limitations. SGN does not include avoidance behavior with respect to entire groups. Furthermore, SGN does not give the user control over the temporary splitting behavior of a group. Groups may split and re-establish coherence, but the splitting phase treats the group members as individuals.

In addition, the vision of an agent is still an approximation and is not influenced by the environment. For instance, agents in open spaces should see hundreds of meters ahead, while their vision should be limited in indoor environments. The computational complexity of maintaining actual vision is a bottleneck that justifies using approximated vision as used in previous methods, e.g. by [Moussaïd et al. 2011].

7 Conclusion and Future Work

In this paper, we have proposed a new method named *Social Groups and Navigation* (SGN) to simulate social-group behavior for virtual pedestrians. SGN is based on the vision-based collision-avoidance method by [Moussaïd et al. 2011] and the social-force model by [Moussaïd et al. 2010]. It is designed to let small social groups stay in coherent and socially-friendly formations without explicitly modeling such formations. Furthermore, SGN lets a group re-organize itself when coherence is lost during the simulation. Thus, SGN handles both global and local aspects of the path planning pipeline to ensure coherent and socially-friendly group navigation. We have shown experimentally that our method improves over existing methods by generating more coherent and socially-friendly walking behavior. Furthermore, our method runs at interactive rates for large numbers of agents, and it can be integrated into any existing crowd simulation framework that supports global path planning, local path following and micro-behavior in separate steps [Curtis et al. 2014; van Toll et al. 2015]. SGN is thus well-suited for gaming and simulation applications that require believable and efficient group behavior for virtual pedestrians.

References

- COLEMAN, J. S., AND JAMES, J. 1961. The equilibrium size distribution of freely-forming groups. *Sociometry* 24, 1, pp. 36–45.
- COSTA, M. 2010. Interpersonal distances in group walking. *Journal of Nonverbal Behavior* 34, 1, 15–26.
- CURTIS, S., BEST, A., AND MANOCHA, D. 2014. Menge: A modular framework for simulating crowd movement. Tech. rep., University of North Carolina at Chapel Hill.
- FIORINI, P., AND SHILLER, Z. 1998. Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research* 17, 7, 760–772.
- FRIDMAN, N., KAMINKA, G. A., AND ZILKA, A. 2013. The impact of culture on crowd dynamics: An empirical approach. In *Proc. International Conference on Autonomous Agents and Multi-agent Systems*, International Foundation for Autonomous Agents and Multiagent Systems, AAMAS '13, 143–150.
- FRUIN, J. J. 1971. *Pedestrian Planning and Design*. Metropolitan Association of Urban Designers and Environmental Planners.
- GERAERTS, R. 2010. Planning short paths with clearance using Explicit Corridors. In *Proc. IEEE International Conference on Robotics and Automation*, 1997–2004.
- HUANG, T., KAPADIA, M., BADLER, N., AND KALLMANN, M. 2014. Path planning for coherent and persistent groups. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, 1652–1659.
- JAKLIN, N., COOK IV, A., AND GERAERTS, R. 2013. Real-time path planning in heterogeneous environments. *Computer Animation and Virtual Worlds* 24, 285–295.
- JAMES, J. 1953. The Distribution of Free-Forming Small Group Size. *American Sociological Review* 18, 5, 569–570.
- KALLMANN, M. 2010. Shortest paths with arbitrary clearance from navigation meshes. In *Proc. Eurographics / SIGGRAPH Symposium on Computer Animation (SCA)*.
- KAMPHUIS, A., AND OVERMARS, M. H. 2004. Finding paths for coherent groups using clearance. In *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SCA '04, 19–28.
- KARAMOUZAS, I., AND OVERMARS, M. 2012. Simulating and evaluating the local behavior of small pedestrian groups. *IEEE Transactions on Visualization and Computer Graphics* 18, 3, 394–406.
- KARAMOUZAS, I., GERAERTS, R., AND OVERMARS, M. 2009. Indicative routes for path planning and crowd simulation. *4th International Conference on Foundations of Digital Games*, 113–120.
- KIMMEL, A., DOBSON, A., AND BEKRIS, K. 2012. Maintaining team coherence under the velocity obstacle framework. In *Proc. 11th International Conference on Autonomous Agents and Multiagent Systems - Volume I*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, AAMAS '12, 247–256.
- KÖSTER, G., TREML, F., SEITZ, M., AND KLEIN, W. 2014. Validation of crowd models including social groups. In *Pedestrian and Evacuation Dynamics 2012*, U. Weidmann, U. Kirsch, and M. Schreckenberg, Eds. Springer International Publishing, 1051–1063.
- LIDDLE, J., SEYFRIED, A., STEFFEN, B., KLINGSCH, W., RUPPRECHT, T., WINKENS, A., AND BOLTES, M. 2011. Microscopic insights into pedestrian motion through a bottleneck, resolving spatial and temporal variations. *arXiv:1105.1532v1*.
- MOUSSAÏD, M., PEROZO, N., GARNIER, S., HELBING, D., AND THERAULAZ, G. 2010. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS ONE* 5, 4, e10047.
- MOUSSAÏD, M., HELBING, D., AND THERAULAZ, G. 2011. How simple rules determine pedestrian behavior and crowd disasters. *Proc. National Academy of Sciences* 108, 17, 6884–6888.
- MUSSE, S. R., AND THALMANN, D. 1997. A model of human crowd behavior : Group inter-relationship and collision detection analysis. In *Computer Animation and Simulation 97*, D. Thalmann and M. Panne, Eds., Eurographics. Springer Vienna, 39–51.
- OLIVA, R., AND PELECHANO, N. 2015. Clearance for diversity of agents sizes in navigation meshes. *Computers & Graphics* 47, 0, 48 – 58.
- ONDŘEJ, J., PETTRÉ, J., OLIVIER, A.-H., AND DONIKIAN, S. 2010. A synthetic-vision based steering approach for crowd simulation. In *ACM SIGGRAPH 2010 Papers*, ACM, New York, NY, USA, SIGGRAPH '10, 1–9.
- PARK, S. I., QUEK, F., AND CAO, Y. 2012. Modeling social groups in crowds using common ground theory. In *Proc. Winter Simulation Conference*, Winter Simulation Conference, WSC '12, 113:1–113:12.
- PELECHANO, N., ALLBECK, J., AND BADLER, N. 2008. *Virtual Crowds: Methods, Simulation, and Control (Synthesis Lectures on Computer Graphics and Animation)*. Morgan and Claypool Publishers.
- QIU, F., AND HU, X. 2010. Modeling group structures in pedestrian crowd simulation. *Simulation Modelling Practice and Theory* 18, 2, 190–205.
- THALMANN, D., AND MUSSE, S. R. 2013. *Crowd Simulation, Second Edition*. Springer.
- VAN DEN BERG, J., GUY, S. J., LIN, M. C., AND MANOCHA, D. 2009. Reciprocal n-body collision avoidance. In *International Symposium of Robotics Research*.
- VAN TOLL, W., JAKLIN, N., AND GERAERTS, R. 2015. Towards believable crowds: A generic multi-level framework for agent navigation. In *ASCI.OPEN 2015*.
- WEIDMANN, U. 1992. *Transporttechnik der fussgänger. IVT, Institut für Verkehrsplanung, Transporttechnik, Strassen-und Eisenbahnbau 90*.
- WU, Q., JI, Q., DU, J., AND LI, X. 2013. Simulating the local behavior of small pedestrian groups using synthetic-vision based steering approach. In *Proc. 12th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry*, ACM, VRCAI '13, 41–50.
- XU, S., AND DUH, H.-L. 2010. A simulation of bonding effects and their impacts on pedestrian dynamics. *Intelligent Transportation Systems, IEEE Transactions on* 11, 1, 153–161.
- ZIPF, G. K. 1949. *Human Behavior and the Principle of Least Effort*. Addison-Wesley Press.




Figure 6: Results from our comparison of SGN with Moussaïd et al. in the first four of our scenarios. We show the average group coherence, partial sociality, and total sociality (all in %) for both methods and varying group sizes.




Figure 7: Evacuation times for the room-evacuation scenario. Due to spacial constraints, for SGN and Moussaïd et al., we used a fixed starting configuration for all agents with no randomness involved. The results show the evacuation times for a single run of each of the two methods. For comparison, we list the mean evacuation times of the method by [Köster et al. 2014] for individuals and for groups of 2 and 4, and the ground-truth data for individuals as obtained by [Liddle et al. 2011].




Figure 8: The evacuation times for the room-evacuation scenario using our SGN method for different agent radii and group sizes.




Figure 9: Top: The average time per step we achieved with a serial and parallel execution of SGN for 1980 agents and different group sizes. Bottom: The frame rates we achieved with a serial and parallel execution of SGN for 1980 agents and different group sizes.