Real-Time Path Planning and Navigation for Multi-agent and Crowd Simulations

Ming C. Lin, Avneesh Sud*, Jur Van den Berg,Russell Gayle, Sean Curtis, Hengchin Yeh, Stephen Guy, Eric Andersen**, Sachin Patil, Jason Sewall, and Dinesh Manocha

Department of Computer Science
University of North Carolina
Chapel Hill, NC, 27599-3175 U.S.A.
{lin,sud,berg,rgayle,seanc,yero,sjguy,
andersen,sachin,sewall,dm}@cs.unc.edu
http://gamma.cs.unc.edu/

Abstract. We survey some of our recent work on real-time path planning and navigation of multiple autonomous agents in dynamic environments. The driving application of our work are real-time crowd simulation for computer games, virtual environments, and avatar-based online 3D social networks. We also present extensions to these methods for accelerating the overall simulation and for modeling more complex behaviors. Finally, we present some preliminary results from our simulations.

Keywords: Velocity Obstacles, Agent-based Simulation, Roadmaps.

1 Introduction

Modeling of crowds, multiple agents and swarm-like behaviors has been widely studied in computer graphics, robotics, architecture, physics, psychology, social sciences, and civil and traffic engineering. Swarms and crowds, ubiquitous in the real world from groups of humans to schools of fish, are vital phenomena to understand and model. In this paper, we address the problems of collision-free path computation for multiple independent agents moving in a dynamic complex game environment and a closely related problem of local navigation among multiple moving agents.

Agent-based techniques focus on modeling individual behaviors and intents. They offer many attractive benefits, as they often result in more realistic and detailed simulations. However, individuals constantly adjust their behavior according to dynamic factors (e.g. another approaching individual) in the environment. Therefore, one of the key challenges in a large-scale agent-based simulation is global collision-free path planning for each virtual agent. The path planning

* Now at Microsoft Coproration.
** Now at University of Washington.

A. Egges, A. Kamphuis, and M. Overmars (Eds.): MIG 2008, LNCS 5277, pp. 23–32, 2008.

problem can become very challenging for real-time applications with a large group of moving agents, as each is a dynamic obstacle for others. Moreover, there may be other dynamic obstacles in the scene, and the underlying applications cannot make any assumptions about their motion. Many prior techniques are either restricted to static environments, or only perform local collision avoidance computations, or may not scale to complex environments with hundreds of agents. The use of solely local methods can result in unnatural behavior or "getting stuck" in local minima. As a result, it is important to combine them with global methods in order to guarantee path convergence for each agent. These problems tend to be more challenging in a dynamically changing scene with multiple moving virtual agents.

We present several complementary approaches for path planning and navigation of multiple virtual agents in a dynamic environment. These methods can be applied to interactive crowd simulation and coordination of multiple autonomous agents in computer games. These include (a) Multi-agent Navigation Graph (MaNG) [37]; (b) Adaptive Elastic ROadmaps (AERO) [38]; and (c) navigation using Reciprocal Velocity Obstacles (RVO) [44]. Each of these methods are suitable for various environments of different characteristics. We also demonstrate how these algorithms can be accelerated by "Pedestrian Levels of Detail" (PLODs) and extended to capture complex human behaviors using "Composite Agents" (CA).

2 Related Work

In this section, we give a brief overview of prior work in this area.

Multiple Moving Entities: The problem of motion planning among multiple agents moving simultaneously in an environment is challenging because of the addition of the number of degrees of freedom which becomes large. The problem was proved to be intractable [18]. It is a specific case of the general time-varying motion planning problem. Two key approaches are known to address this problem: centralized and decoupled planning.

The *centralized* approaches [17,18] consider the sum of all the robots as a single one. For that, configuration spaces of individual robots are combined (using Cartesian product) in a composite one, in which a solution is searched. As the dimension of the composite space grows with the number of degrees of freedom added by each entity, the problem complexity becomes prohibitively high.

Contrarily, the *decoupled* planners proceed in a distributed manner and coordination is often handled by exploring a *coordination space*, which represents the parameters along each specific robot path. Decoupled approaches [34,45] are much faster than centralized methods, but may not be able to guarantee completeness.

Dynamic Environments: Evolving elements significantly increase the difficulty of the motion planning problem. In fact, motion planning for a single disc with bounded velocity among rotating obstacles is PSPACE-hard [18].

However, there have been many attempts to provide practical methods to cope with changing environments. For example, Stentz et al. proposed the D* deterministic planning algorithm to repair previous solutions instead of re-planning from scratch [15,35].

There are two main approaches for adapting randomized planners to dynamic environments [10,19]. The first one includes both PRMs and RRTs that reuse previously computed information to aid in finding a new path [5,11,12,20]. The second integrates obstacle motion directly into the planning process. Some variations plan directly in a C-space augmented with a time parameter [28].

Rather than changing the roadmap, other work on dynamic environments has focused on adjusting or modifying the path. Potential field planners use gradient descent to move toward a goal, at a potential sink [14]. Building on these ideas, several variation of dynamic changing or elastic roadmaps have been proposed [8,29,46].

Agent Simulation and Crowd Dynamics: Modeling of collective behaviors has been heavily studied in many fields [9,13,23,32,36]. Simulation of multiple avatars agents and crowds have also been well studied in graphics and VR [1,30,33,41,42]. They differ based on problem decomposition (discrete vs continuous), stochastic vs deterministic, etc.

Numerous approaches have been proposed using variants of agent-based methods, rule-based techniques, and social force models [4,22,24,25,26,31,39,43]. They focus mainly on local planning of agents, and can exhibit emergent behaviors. Global methods using path planning algorithms have also been suggested [2,6,16,27,33,40], but are mostly limited to static environments. More recently algorithms have been proposed to extend the roadmap-based methods to dynamic environments and multiple agents [7,8,21,27,47] in relatively simple environments with only a few entities. Hybrid methods that treat each agent as a particle in a dynamical system combined with high-level planning can lead to instability in the simulation or entrapment at local minima in the energy potential (See examples discussed in [42]).

3 Multi-agent Navigation Graphs

We introduce a new data structure called "multi-agent navigation graph" or MaNG and compute it efficiently using GPU-accelerated discrete Voronoi diagrams. Voronoi diagrams have been widely used for path planning computations in static environments and we extend these approaches to dynamic environments. Voronoi diagrams encode the connectivity of the free space and provide a path of maximal clearance for an agent from other obstacles. In order to use them for multiple moving agents in a dynamic scene, prior approaches compute the Voronoi diagram for each agent separately, treating the other agents and the environment as obstacles. This approach can become very costly for large number of virtual agents. Instead, we compute the *second order* Voronoi diagram of all the obstacles and agents, and show that the second order Voronoi diagram provides *pairwise* proximity information for all the agents simultaneously.

Therefore, we combine the first and second order Voronoi graphs to compute a single MaNG which provides global path planning of multiple virtual agents.

We describe novel algorithm for computing the first and second order diagrams of agents and obstacles using graphics hardware. In practice, our approach can handle environments with few hundred agents and obstacles in real time. We make no assumption on the motion, and the computed path lies along the Voronoi boundary, i.e. farthest from the obstacles. We also observe many interesting and emergent behaviors based on our method [37].

4 Heterogeneous Crowd Simulations Using AERO

We present a new algorithm for real-time simulation of large-scale heterogeneous crowds in complex dynamic environments. In his pioneering work, Gustave Le Bon [3] defined a *heterogeneous crowd* as consisting of many dissimilar types of groups, each with potentially independent behavior characteristics and goals. Examples include large exposition halls, wide festival arenas, busy urban street scenes, etc. Real-time simulation of heterogeneous crowds is highly challenging because pedestrian dynamics exhibits a rich variety of both independent and collective effects, such as lane formations, oscillations at bottlenecks, chemotaxis and panic effects.

Our approach is based on a novel concept called *"Adaptive Elastic ROadmaps" (AERO)* [38]. It is a connectivity graph structure that is lazily computed using a generalized crowd dynamics model. Specifically, we use a reactive roadmap computation algorithm [8] that updates the links of the roadmap graphs using a particle-based dynamics simulator. This algorithm includes path modification, addition and deletion of links, as well as multiple lane formations around each link. Our dynamics model simultaneously captures macroscopic mutual interaction among multiple moving *groups* and microscopic local forces among *individual* pedestrians or agents. We use AERO to perform dynamic, global path planning based on this force model. This approach has been successfully demonstrated in complex urban scenes, trade-shows as well as terrain scenarios with multiple moving vehicles.

5 Navigation Using Reciprocal Velocity Obstacles

We also present another technique for local navigation. Our approach is decomposed into two levels. The higher-level deals with the global path planning towards the goal using a precomputed roadmap [18], and the lower-level addresses local collision avoidance and navigation using Reciprocal Velocity Obstacles (RVO) [44]. We assume that the other agents are dynamic obstacles whose future motions are predicted as linear extrapolations of their current velocities. RVO provides a principle to select a velocity for agent A_i and implicitly assumes that the other agents A_j use similar collision avoidance reasoning. Informally speaking, this means that agent A_i does only half of the effort to avoid a collision with agent A_j, and assumes that the other agent will take care of the other half.

The Reciprocal Velocity Obstacle $RVO_j^i(\mathbf{v}_j, \mathbf{v}_i)$ of agent A_j to agent A_i is defined as the set consisting of all those velocities $\mathbf{v}_i$ for A_i that will result in a collision at some moment in time with agent A_j, *if* agent A_j chooses a velocity in its Reciprocal Velocity Obstacle as well. n other words, if A_i chooses its velocity outside the Reciprocal Velocity Obstacle, A_i is guaranteed to avoid collisions with agent A_j, *provided* that A_j applies the same navigation algorithm, and chooses a new velocity outside its Reciprocal Velocity Obstacle. As there are multiple agents around, each agent computes its Reciprocal Velocity Obstacle as the *union* of the Reciprocal Velocity Obstacles corresponding to the individual agents, and selects a velocity outside this union to avoid collisions. Also, the Reciprocal Velocity Obstacle method is guaranteed to avoid oscillatory behavior of the agents. This basic algorithmic framework can also incorporate kinodynamic constraints, orientation, and visibility regions of each agents, etc.

6 Extensions

In this section, we present techniques to extend these methods for capturing more complex behaviors and accelerate the overall simulations.

6.1 Modeling Complex Behaviors

We introduce the notion of composite agents to effectively model different avatar behaviors for agent-based crowd simulation. Each composite agent consists of a basic agent that is associated with one or more *proxy agents*. A proxy agent possesses the same kind of external properties as an agent, i.e. if every agent's external property consists of the velocity, then the proxy agent must have its own velocity as well. The values of these external properties, can be different, e.g. the proxy can possesses a velocity that is different from anyone else. The external properties of a proxy P_j is denoted as ε_j. The internal state ι_j of a proxy agent, however, need not be the same set of properties of the basic agent. We also define that P_j has access to the internal state ι_i of its parent A_i. We denote the set of all proxy agents being simulated as $Proxies = \bigcup_i proxy(A_i)$

This formulation allows an agent with given physical properties to exercise its influence over other agents and the environment. Composite agents can be added to most agent-based simulation systems and used to model emergent behaviors among individuals. These behaviors include aggression, impatience, intimidation, leadership, trailblazing and approach-avoidance conflict, etc. in complex scenes. In practice, there is negligible overhead of introducing composite agents in the simulation.

6.2 Accelerating Simulations

We can accelerate crowd simulation by using "Pedestrian Levels of Detail" (PLODs). PLODs is a new hierarchical data structure based on Kd-tree

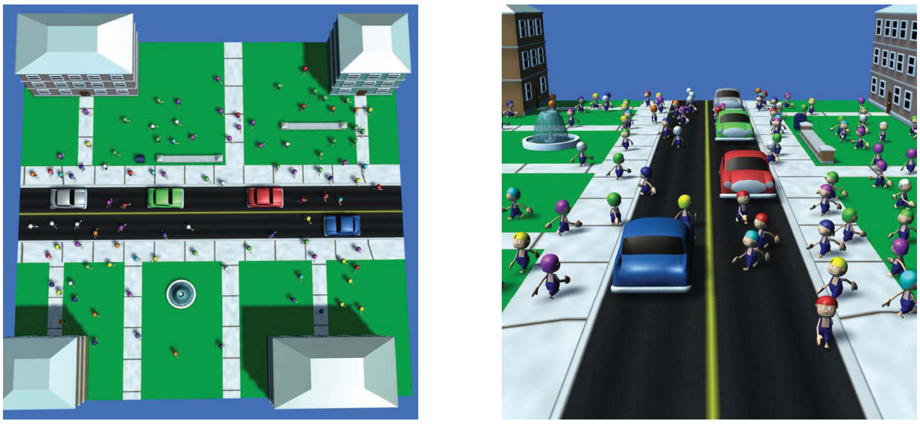

Fig. 1. Multi-agent simulation using MaNG

Fig. 2. Real-time heterogeneous crowd simulation using AERO

selectively and dynamically computed. It is used to adaptively cluster agents to accelerate large-scale simulation of heterogeneous crowds. Our formulation is based on the observation and empirical validation in traffic engineering that crowds exhibit self-organization in pedestrian flows. Depending on their dynamic states (e.g. walking speed, heading directions), spatial proximity, and behavior characteristics, agents in heterogeneous crowds are adaptively grouped and subdivided into PLODs, thus reducing the overall computation of heterogeneous crowd simulation.

7 Results

We demonstrate our approaches on several complex indoor and outdoor scenarios. In the first scenario, we show that we can compute MaNG for real-time path planning of $100 - 200$ agents in real time (approx. 10 fps) in a small town environment (see Fig. 1). In the second scenario shown in Fig. 2, we illustrate the

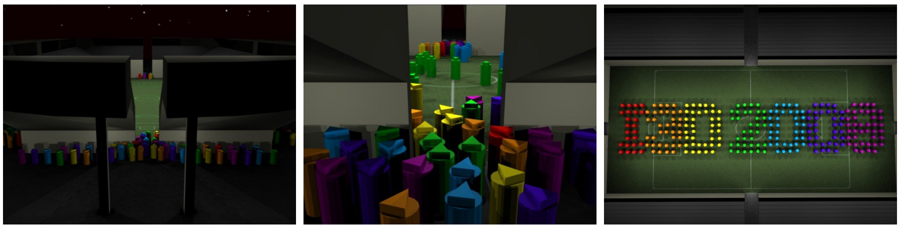

Fig. 3. Real-time Navigation using RVO. 250 agents form the word 'I3D 2008' on the field after entering the stadium; congestion develops near the stadium entrances.

performance of our real-time heterogeneous crowd simulation using AERO on a exhibition Hall scene, consisting of 1,000 pedestrians and 500 booths. Fig. 3 shows pedestrians crossing a street at four corners of crosswalks, demonstrating "lane formation" that naturally emerges using RVO. Fig. 4 demonstrates a variety of behaviors generated using our novel framework of *Composite Agents* on top of any agent-based simulation.

Each one of these methods have its advantages and limitations. The choice of methods depends to some degree on the complexity of the simulation environments, the density of agents and their physical interactions, the diversity of their behaviors, etc.

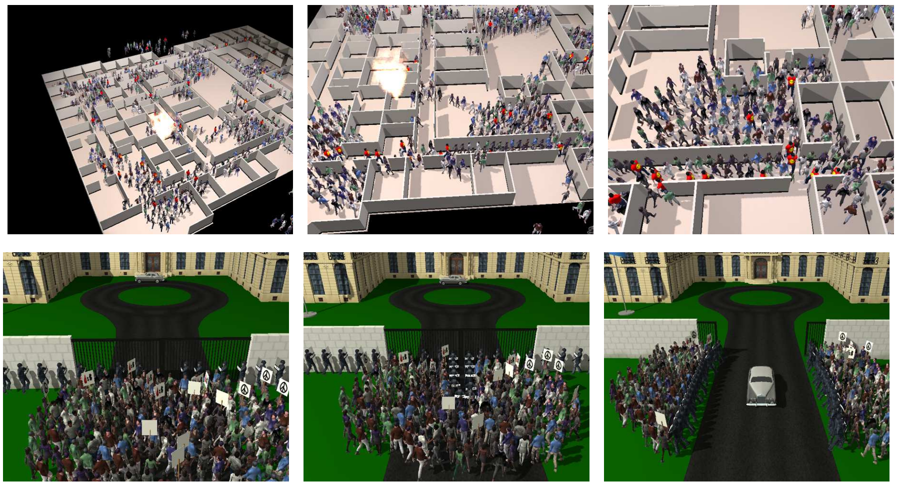

Fig. 4. Modeling of Complex Behaviors. Top row depicts an emergency evacuation in an office building. The agents in red are aggressive agents. They are able to carve their own way through the crowd and exit more quickly than the others. The bottom row simulates a crowd of protesters outside an embassy. Two files of policemen clear the protesters off the road. Notice that when forcing their way into the crowd, even if the actual gap between the individual policemen is enough for a protest or to pass through, the perceived continuity of authority prevents the protesters from breaking the barricade.

8 Discussion

Among these methods, MaNG is better suited for a smaller scene with hundreds of autonomous agents running on a desktop PC with a top-of-the-line GPU. MaNG works well for all types of environments, including open outdoor terrains.

AERO runs at interactive rates for an environment with upto thousands of agents. Some performance gain can be achieved by further exploiting GPUs to compute the interaction forces among agents. AERO currently uses a physics-based social force model. However, a more geometric, velocity-based local navigation method, such as RVO, can be used instead. AERO performs very well for structured urban scenes and modestly well for open landscapes; while the accompanying acceleration method, PLOD, achieves most gain on structured landscapes.

Our current planner that uses RVO as a low-level collision avoidance module is built upon a static, precomputed roadmap. This planning algorithm works best for structured environments. The algorithm scales well as the number of cores increases, thus highly parallelizable to run on many-core architectures, similarly for composite agents.

Acknowledgments. The research presented here is supported in part by Army Research Office, Intel, National Science Foundation, and DARPA/RDECOM.

References

1. Ashida, K., Lee, S.J., Allbeck, J., Sun, H., Badler, N., Metaxas, D.: Pedestrians: Creating agent behaviors through statistical analysis of observation data. In: Proc. Computer Animation (2001)
2. Bayazit, O.B., Lien, J.-M., Amato, N.M.: Better group behaviors in complex environments with global roadmaps. In: Int. Conf. on the Sim. and Syn. of Living Sys. (Alife), pp. 362–370 (2002)
3. Le Bon, G.: The Crowd: A Study of the Popular Mind. Dover Publications (1895)
4. Cordeiro, O.C., Braun, A., Silveria, C.B., Musse, S.R., Cavalheiro, G.G.: Concurrency on social forces simulation model. In: First International Workshop on Crowd Simulation (2005)
5. Ferguson, D., Kalra, N., Stentz, A.: Replanning with RRTs. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (May 2006)
6. Funge, J., Tu, X., Terzopoulos, D.: Cognitive modeling: Knowledge, reasoning and planning for intelligent characters. In: Proc. of ACM SIGGRAPH, pp. 29–38 (1999)
7. Garaerts, R., Overmars, M.H.: The corridor map method: Real-time high-quality path planning. In: ICRA, Roma, Italy, pp. 1023–1028 (2007)
8. Gayle, R., Sud, A., Lin, M., Manocha, D.: Reactive deformation roadmaps: Motion planning of multiple robots in dynamic environments. In: Proc IEEE International Conference on Intelligent Robots and Systems (2007)
9. Helbing, D., Buzna, L., Werner, T.: Self-organized pedestrian crowd dynamics and design solutions. Traffic Forum 12 (2003)
10. Hsu, D., Kindel, R., Latombe, J.-C., Rock, S.: Randomized kinodynamic motion planning with moving obstacles. International Journal of Robotics Research (2002)

11. Jaillet, L., Simeon, T.: A PRM-based motion planning for dynamically changing environments. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (2004)
12. Kallmann, M., Mataric, M.: Motion planning using dynamic roadmaps. In: Proceedings of the IEEE Conference on Robotics and Automation (ICRA) (April 2004)
13. Kamphuis, A., Overmars, M.: Finding paths for coherent groups using clearance. In: Proc. of ACM SIGGRAPH / Eurographics Symposium on Computer Animation, pp. 19–28 (2004)
14. Khatib, O.: Real-time obstable avoidance for manipulators and mobile robots. IJRR 5(1), 90–98 (1986)
15. Koenig, S., Likhachev, M.: Improved fast replanning for robot navigation in unknown terrain. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (May 2002)
16. Lamarche, F., Donikian, S.: Crowd of virtual humans: a new approach for real time navigation in complex and structured environments. Computer Graphics Forum 23(3), 509–518 (2004)
17. Latombe, J.C.: Robot Motion Planning. Kluwer Academic Publishers, Dordrecht (1991)
18. LaValle, S.M.: Planning Algorithms. Cambridge University Press, Cambridge (2006), `http://msl.cs.uiuc.edu/planning/`
19. La Valle, S., Kuffner, J.: Randomized kinodynamic planning. International Journal of Robotics Research (2001)
20. Leven, P., Hutchinson, S.: Toward real-time path planning in changing environments. In: Proceedings of the fourth International Workshop on the Algorithmic Foundations of Robotics (WAFR) (2000)
21. Li, Y., Gupta, K.: Motion planning of multiple agents in virtual environments on parallel architectures. In: ICRA, Roma, Italy, pp. 1009–1014 (2007)
22. Loscos, C., Marchal, D., Meyer, A.: Intuitive crowd behaviour in dense urban environments using local laws. In: Theory and Practice of Computer Graphics (TPCG 2003), pp. 122–129 (2003)
23. MASSIVE (2006), `http://www.massivesoftware.com`
24. Musse, S.R., Thalmann, D.: A model of human crowd behavior: Group inter-relationship and collision detection analysis. In: Computer Animation and Simulation, pp. 39–51 (1997)
25. Pelechano, N., Allbeck, J., Badler, N.: Controlling individual agents in high-density crowd simulation. In: Proc. of ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA) (2007)
26. Pelechano, N., O'Brien, K., Silverman, B., Badler, N.: Crowd simulation incorporating agent psychological models, roles and communication. In: First International Workshop on Crowd Simulation (2005)
27. Pettre, J., Laumond, J.-P., Thalmann, D.: A navigation graph for real-time crowd animation on multilayered and uneven terrain. In: First International Workshop on Crowd Simulation (2005)
28. Petty, S., Fraichard, T.: Safe motion planning in dynamic environments. In: Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, pp. 3726–3731 (2005)
29. Quinlan, S., Khatib, O.: Elastic bands: Connecting path planning and control. In: Proc. of IEEE Conf. on Robotics and Automation (1993)
30. Reynolds, C.: Big fast crowds on ps3. In: Sandbox 2006: Proceedings of the 2006 ACM SIGGRAPH symposium on Videogames, pp. 113–121. ACM Press, New York (2006)

31. Reynolds, C.W.: Flocks, herds, and schools: A distributed behavioral model. In: SIGGRAPH 1987 Proceedings, vol. 21, pp. 25–34 (1987)
32. Schreckenberg, M., Sharma, S.D.: Pedestrian and Evacuation Dynamics. Springer, Heidelberg (2001)
33. Shao, W., Terzopoulos, D.: Autonomous pedestrians. In: SCA 2005: Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation, pp. 19–28. ACM Press, New York (2005)
34. Simeon, T., Leroy, S., Laumond, J.: Path coordination for multiple mobile robots: a geometric algorithm. In: Proc. of IJCAI (1999)
35. Stentz, A.: The focussed D* algorithm for real-time replanning. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) (1995)
36. Still, G.: Crowd Dynamics. PhD thesis, University of Warwik, UK, Ph.D. Thesis (2000)
37. Sud, A., Andersen, E., Curtis, S., Lin, M., Manocha, D.: Real-time path planning for virtual agents in dynamic environments. In: Proc. of IEEE VR, pp. 91–98 (2007)
38. Sud, A., Gayle, R., Andersen, E., Guy, S., Lin, M., Manocha, D.: Real-time navigation of independent agents using adaptive roadmaps. In: Proc. of ACM VRST (2007)
39. Sung, M., Gleicher, M., Chenney, S.: Scalable behaviors for crowd simulation. Computer Graphics Forum 3, 519–528 (2004)
40. Sung, M., Kovar, L., Gleicher, M.: Fast and accurate goal-directed motion synthesis for crowds. In: Proc. of SCA 2005, pp. 291–300 (2005)
41. Thalmann, D., O'Sullivan, C., Ciechomski, P., Dobbyn, S.: Populating Virtual Environments with Crowds. In: Eurographics 2006 Tutorial Notes (2006)
42. Treuille, A., Cooper, S., Popovic, Z.: Continuum crowds. In: Proc. of ACM SIGGRAPH, pp. 1160–1168 (2006)
43. Tu, X., Terzopoulos, D.: Artificial fishes: Physics, locomotion, perception, behavior. In: Glassner, A. (ed.) Proceedings of SIGGRAPH 1994, pp. 43–50 (1994)
44. van den Berg, J., Patil, S., Seawall, J., Manocha, D., Lin, M.: Interactive navigation of individual agents in crowded environments. In: Proc. of ACM Symposium on Interactive 3D Graphics and Games (2008)
45. Warren, C.W.: Multiple path coordination using artificial potential fields. In: Proc. of IEEE Conf. on Robotics and Automation, pp. 500–505 (1990)
46. Yang, Y., Brock, O.: Elastic roadmaps: Globally task-consistent motion for autonomous mobile manipulation. In: Proceedings of Robotics: Science and Systems (August 2006)
47. Zucker, M., Kuffner, J., Branicky, M.: Multipartite rrts for rapid replanning in dynamic environments. In: Proc. IEEE Int. Conf. on Robotics and Automation (2007)