

## SPECIAL ISSUE PAPER

**Footstep navigation for dynamic crowds**

Shawn Singh\*, Mubbasir Kapadia, Glenn Reinman and Petros Faloutsos

Department of Computer Science, University of California, Los Angeles, USA

**ABSTRACT**

The majority of steering algorithms output only a force or velocity vector to an animation system, without modeling the constraints and capabilities of human-like movement. This simplistic approach lacks control over how a character should navigate. This paper proposes a steering method that uses *footsteps* to navigate characters in dynamic crowds. Instead of an oriented particle with a single collision radius, we model a character's center of mass and footsteps using a 2D approximation of an inverted spherical pendulum model of bipedal locomotion. We use this model to generate a timed sequence of footsteps that existing animation techniques can follow exactly. Our approach not only constrains characters to navigate with realistic steps but also enables characters to intelligently control subtle *navigation* behaviors that are possible with exact footsteps, such as side-stepping. Our approach can navigate crowds of hundreds of individual characters with collision-free, natural steering decisions in real-time. Copyright © 2011 John Wiley & Sons, Ltd.

**KEYWORDS**

crowds; footprints; footsteps; navigation; steering

**\*Correspondence**

Shawn Singh, Department of Computer Science, University of California, Los Angeles, Boelter Hall 4531F, Los Angeles, CA 90095, USA.

E-mail: shawnsin@cs.ucla.edu

**1. INTRODUCTION**

The majority of previous steering algorithms represent a character as an oriented particle that moves by choosing a force or velocity vector. Often, orientation is heuristically chosen to be the particle's velocity. This approach has the two key disadvantages:

- *Limited locomotion constraints*: Most steering algorithms do not account for constraints of real human movement. Trajectories may have discontinuous velocities, oscillations, awkward orientations, or may try to move a character unnaturally, and these side-effects make it harder to animate the character intelligently.
- *Limited navigation control*: It is common to assume that an animation system will know how to interpret a vector-based steering decision. In practice, a vector does not have enough information to indicate appropriate maneuvers, such as side-stepping versus reorienting the torso, stepping backwards versus turning around, planting a foot to change momentum quickly, or carefully placing steps in exact locations.

We propose to generate sequences of footsteps as the output of navigation. Since there are already several animation techniques that can animate a character to follow

timed footsteps exactly [1–6], the main challenge and focus of our work is how to *generate* footsteps as the output of navigation. Footsteps are an intuitive abstraction for most locomotion tasks, and they provide precise, unambiguous spatial and timing information to animation.

In our system, each step is defined by a 2D parabolic trajectory that approximates the motion of a 3D inverted pendulum. The location, orientation, and timing of footsteps are derived from these trajectories. We use a best-first search to plan a sequence of space-time parabolic trajectories and the associated footsteps that avoids time-varying collisions, satisfies footstep constraints for natural locomotion, and minimizes the effort to reach a local goal. Characters successfully avoid collisions with each other and choose steps that correspond to natural and fluid motion, including precise timing. Because the most significant biomechanics constraints are already taken into account in our model, integrating our results with an existing animation algorithm that follows footsteps is straightforward and results in navigation that is often richer and less awkward than vector-based navigation.

*Contributions.* This paper presents a new approach to steering in dynamic crowds that uses a simple biomechanically-based footstep model combined with space-time

planning. Our work demonstrates that a steering algorithm can have better navigation features than a vector interface, while still retaining fast performance. These features include: short-term space–time planning, dynamic collision bounds, appropriate movement constraints, and more precise navigation control. Because substantial work already exists to animate characters to follow exact footsteps including timing information, we focus on the navigation: how to generate biomechanically plausible footsteps for dynamic crowds.

## 2. RELATED WORK

Two widely accepted strategies are (1) the social forces model [7], which associates a small force field around agents and obstacles, and (2) the steering behaviors model [8], where forces are procedurally computed to perform desired functions such as seek, flee, pursuit, evasion, and collision avoidance. Many works are extensions or elaborations of these two ideas [9–17]. A more complete survey of collision avoidance, navigation, and crowd simulation work can be found in Reference [18]. The common theme in these works is the use of force or velocity vectors as navigation decisions, which has the limitations described above.

Only a few steering techniques take into account locomotion constraints that an animation system will have. Paris and Donikian [19] demonstrate a framework where the animation module can potentially tell steering that an action is not plausible. Musse and Thalmann [20] and Shao and Terzopoulos [21] both address higher-level aspects of pedestrians, and their navigation modules output a choice from a set of navigation behaviors that correspond directly to animations the character can produce. Van Basten and Egges [22] discuss problems of interfacing navigation with animation, proposing abstractions that reduce such discrepancies.

Another approach to navigation is to plan sequences of motion clips [23], demonstrated this is possible in real-time for crowds, by precomputing a tree of all possible sequences of motion clips. However, a large number of motion clips would be needed to emulate the versatility of far fewer stepping options. The technique of precomputing a search tree can also be applied to our footstep planner, but our approach is scalable even without this extension.

*Footsteps.* Several animation techniques, academic and commercial, can *follow* a given sequence of footsteps [1–6], and more. Animation methods in these works include forward and inverse kinematics, physically based control, and motion capture.

The challenge of *generating* footsteps has so far only been explored for single characters in static environments. Research in robotics [24–28] explores autonomous footplacement to avoid obstacles while navigating towards a goal. Their focus is practical robot control, and so they do not consider issues of real human locomotion. Torkos and Van de Panne [1] generate footsteps to randomly wander,

changing direction if nearby objects are too close, used to demonstrate their animation system. Chung and Hahn [2] input a trajectory, and generate footsteps by aligning each step to the orientation of the trajectory, with smaller footsteps around curves. Choi *et al.* [29] use roadmaps to plan sequences of steps, choosing from steps that are possible with the given motion clips and requiring costly roadmap construction and footstep verification. Zhang *et al.* [30] propose a hierarchical planning approach that computes full-body motion including footsteps for tasks in highly constrained environments. Recently several papers have considered footsteps as a way of guiding controllers for physically-based character animations [3,31].

## 3. FOOTSTEP MODEL

The primary data structure in our model is a *footstep*, which includes: (1) the position, velocity, and timing of the character’s center of mass trajectory, (2) the location and orientation of the foot, and (3) the cost of taking the step. In this section, we describe these aspects of a footstep, as well as the constraints for choosing footsteps.

*Center of Mass Trajectory.* The analogy between human locomotion and the inverted pendulum is well known [32]; the pendulum pivot represents a point on or near a footstep, while the pendulum mass represents a character’s center of mass. We define a 2D analytical approximation to the dynamics of an inverted spherical pendulum using parabolas. Piecewise parabolic curves are enough to capture the variety of trajectories that a human’s center of mass will have: varying curvature, speed, and step sizes. Each step is a parabola defined with the following parameters in local space:

$$(x(t), y(t), \dot{x}(t), \dot{y}(t)) = (v_{x_0}t, \alpha t^2, v_{x_0}, 2\alpha t) \quad (1)$$

such that both  $v_{x_0}$  and  $\alpha$  are positive.

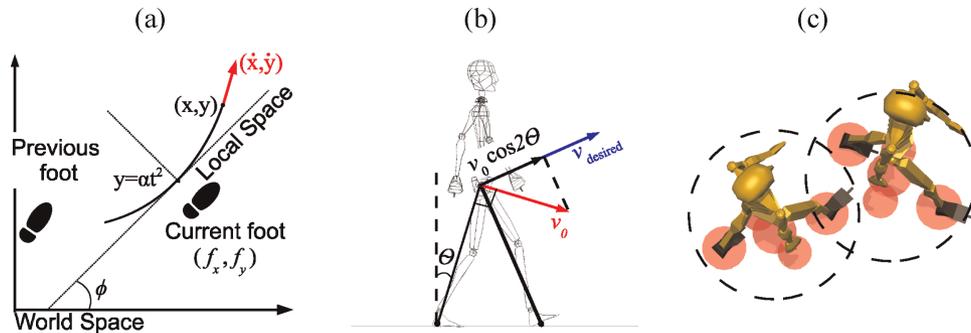
Equation 1 allows us to *analytically* evaluate the position and velocity of a character’s center of mass at any time  $t$ . This makes it practical to search through many possible trajectories for many characters in real-time.

### 3.1. Footstep Actions

The state of the character  $s \in \mathcal{S}$  is defined as follows (Figure 1a):

$$s = \{(x, y), (\dot{x}, \dot{y}), (f_x, f_y), f_\phi, I \in \{\text{L}, \text{R}\}\}$$

where  $(x, y)$  and  $(\dot{x}, \dot{y})$  are the position and velocity of the center of mass of the character at the *end* of the step,  $(f_x, f_y)$  and  $f_\phi$  are the location and orientation of the foot, and  $I$  is an indicator of which foot (left or right) is taking the step. The state space  $\mathcal{S}$  is the set of valid states that satisfy the constraints described below.



**Figure 1.** Our footstep model. (a) Depiction of state and action parameters. (b) A sagittal view of the pendulum model used to estimate energy costs. (c) The collision model uses five circles that track the torso and feet over time, allowing tighter configurations than a single coarse radius.

A *footstep action* determines the next parabolic trajectory, defined as  $a \in \mathcal{A}$ :

$$a = \{\phi, v_{desired}, T\}$$

where  $\phi$  is the desired orientation of the parabola,  $v_{desired}$  is the desired initial speed of the center of mass, and  $T$  is the desired time duration of the step. The action space  $\mathcal{A}$  is the set of valid footstep actions, where the input and output states are both valid. Note that when the character's previous step is fixed, varying  $\phi$  directly affects the width of the parabolic trajectory, thus allowing a large variety of step choices.

A key aspect of the model is the transition function,  $s' = \text{createFootstep}(s, a)$ . This function receives a desired footstep action  $a$  and a state  $s$  and returns a new state  $s'$  if the action is valid. It is implemented as follows. First,  $\phi$ , which indicates the orientation of the parabola, is used to compute a transform from world space to local parabola space. Then, the direction of velocity  $(\dot{x}, \dot{y})$  from the end of the previous step is transformed into local space, normalized, and re-scaled by the desired speed  $v_{desired}$ . With this local desired velocity, there is enough information to solve for  $\alpha$ , and then Equation 1 is used to compute  $(x, y)$  and  $(\hat{x}, \hat{y})$  at the end of the next step. In local space, the foot location is always located at  $(f_x, f_y) = (0, -d)$ , where  $d$  is describes the distance between a character's foot and center at rest. Finally, all state information is transformed back into world space, which serves as the input to create the next footstep.

## 3.2. Locomotion Constraints

### 3.2.1. Biomechanical Properties.

Several properties of human locomotion are automatically enforced by the definition of our model. The piecewise parabola will be G-1 continuous, and the center of mass will remain between the two feet by enforcing the local-space parabola remains positive. Our footstep model offers a number of intuitive parameters with meaningful defaults and well-defined physical meaning. These parameters

include the height of the character's center of mass, the min, max, and preferred step timing and stride length, the preferred and max velocities of the character's walk, the interval of valid foot orientations, etc. If these constraints are violated, the footstep is considered invalid. A user can modify these parameters to create new locomotion styles. For example, restricting the valid range of step timing and output velocity for one foot results in asymmetric limping, like an injured character.

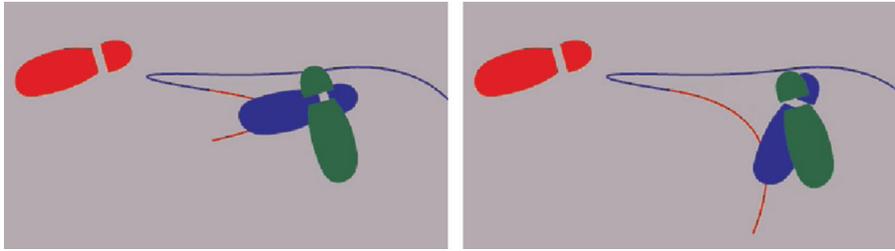
### 3.2.2. Footstep Orientation.

Intuitively, it may seem that footstep orientations must be an additional control parameter when creating a footstep. However, the choice of footstep orientation has no direct effect on the dynamics of the center of mass trajectory; the foot orientation only constrains the options for current trajectory and future footsteps. This is a key aspect to our model's efficiency—instead of increasing the dimensionality of our search space to include foot orientation, we use orientation to constrain the search space of a lower dimensional system.

To implement this constraint, we compute an interval  $[f_{\phi_{inner}}, f_{\phi_{outer}}]$  of valid foot orientations. This interval is constrained by the same interval from the previous step, and further constrained by the parabola orientation  $\phi$  used to create the next footstep (Figure 2):

$$\begin{aligned} & [f_{\phi_{next\_inner}}, f_{\phi_{next\_outer}}] \\ &= \left[ f_{\phi_{prev\_outer}}, f_{\phi_{prev\_inner}} + \frac{\pi}{2} \right] \cap [\phi, \text{atan2}(\dot{y}, \dot{x})] \end{aligned}$$

If this intersection becomes an empty set, that implies that no foot orientation can satisfy the step constraints, so the step is invalid. Note the ordering of bounds in these intervals; the next foot's outer bound is constrained by the previous foot's inner bound. In words, the interval  $[\phi, \text{atan2}(\dot{y}, \dot{x})]$  describes two constraints: (1) the character would not choose a foot orientation that puts his center of mass on the outer side of the foot, (2) a human would rarely orient the next step more outwards than the direction of momentum; violating this constraint would put the



**Figure 2.** An interval of valid foot orientations (the blue and green feet) is maintained for each step, constrained by the previous step (red foot) and the chosen trajectory (red line).

character's center of mass on the wrong side of the foot. The exact orientation is chosen as a fast postprocess, described below.

### 3.2.3. Space-Time Collision Model.

For any given footstep, our model computes the *time-varying* collision bounds of the character at any exact time. To determine if a footstep causes a collision, we iterate over several time-steps within the footstep and query the collision bounds of nearby characters for that time. The collision bounds are five circles, depicted in Figure 1c. Each circle associated with a foot exists while the foot is planted on the ground. The three circles associated with the torso are placed on the center of mass, which moves along the parabola over several time-steps. If any of these circles collide with an obstacle or another character's circles, the footstep is considered invalid.

## 3.3. Cost Function

We define the cost of a given step as the energy spent to execute the footstep action. We model three forms of energy expenditure for a step: (1)  $\Delta E_1$ , a fixed rate of energy the character spends per unit time, (2)  $\Delta E_2$ , the work spent due to ground reaction forces to achieve the desired speed, and (3)  $\Delta E_3$ , the work spent due to ground reaction forces accelerating the center of mass along the trajectory. The total cost of a footstep action transitioning a character from  $s$  to  $s'$  is given by:

$$c(s, s') = \Delta E_1 + \Delta E_2 + \Delta E_3 \quad (2)$$

### 3.3.1. Fixed Energy Rate.

The user defines a fixed rate of energy spent per second, denoted as  $R$ . For each step, this energy rate is multiplied by the time duration of the step  $T$  to compute the cost:

$$\Delta E_1 = R \cdot T \quad (3)$$

This cost is proportional to the amount of time it takes to reach the goal, and thus minimizing this cost corresponds to the character trying to minimize the time it spends walking to his goal. We found that good values for  $R$  are roughly proportional to the character's mass.

### 3.3.2. Ground Reaction Forces.

As a character pushes against the ground, the ground exerts equal and opposite forces on the character. We model three aspects of ground reaction forces that are exerted on the character's center of mass, from the study of biomechanics. The geometry and notation of the cost model is shown in Figure 1b. First, at the beginning of a new step (heel-strike), some of the character's momentum dissipates into the ground. We estimate this as an instantaneous loss of momentum along the pendulum shaft, reducing the character's speed from  $v_0$  to  $v_0 \cos(2\theta)$ . In order to resume a desired speed, the character actively exerts additional work on his center of mass, computed as:

$$\Delta E_2 = \frac{m}{2} |(v_{\text{desired}})^2 - (v_0 \cos(2\theta))^2| \quad (4)$$

This cost measures the effort required to choose a certain speed. At every step, some energy is dissipated into the ground, and if a character wants to maintain a certain speed, it must actively add the same amount of energy back into the system. On the other hand, not all energy dissipates from the system after a step, so if the character wants to come to an immediate stop, the character also requires work to remove energy from the system. Minimizing this cost corresponds to finding footsteps that require less effort, and thus tend to look more natural. Furthermore, when walking with excessively large steps,  $\cos(2\theta)$  becomes smaller, implying that more energy is lost per step.

It should be noted that there is much more complexity to real bipedal locomotion than this cost model. For example, the appropriate bending of knees and ankles and the elasticity of human joints can significantly reduce the energy lost per step, reducing the required work for a real human. While the model is not an accurate measurement of energy spent, it is sufficient for *comparing* the effort of different steps.

$\Delta E_2$  captures only the cost of changing a character's momentum at the beginning of each step. The character's momentum may also change during the trajectory. For relatively straight trajectories, this change in momentum is mostly due to the passive inverted pendulum dynamics that requires no active work. However, for trajectories of high curvature, a character spends additional energy to change his momentum. We model this cost as the work required to

change momentum (denoted as  $P$ ) over the length of the step, weighted by constant  $w$ :

$$\Delta E_3 = w \cdot \frac{dP}{dt} \cdot \text{length} = w \cdot m\alpha \cdot \text{length} \quad (5)$$

Note that  $\alpha$  is the same coefficient in Equation 1, the acceleration of the trajectory.  $\alpha$  increases if the curvature of the parabola is larger, and also if the speed of the character along the trajectory is larger. Minimizing this cost corresponds to preferring straight steps when possible, and preferring to go slower (and consequently, taking smaller steps) when changing the direction of momentum significantly. The weight  $w$  can be adjusted to change whether it costs more energy to walk around an obstacle or to stop and wait for the obstacle to pass. We found good values of  $w$  to be between 0.2 and 0.5, meaning that 20 to 50 per cent of the curvature is due to the character's active effort, and the rest due to the passive inverted pendulum dynamics.

## 4. Generating Sequences of Steps

### 4.1. Discretizing Action Space

The choices for a character's next step are generated by discretizing the action space  $\mathcal{A}$  described above, in all three dimensions and using the `createFootstep( $s, a$ )` function to compute the new state and cost of each action. We have found that  $v_{\text{desired}}$  and  $T$  can be discretized extremely coarsely, as long as there are at least a few different speeds and timings. Further optimizations are made by observing that speed  $v_{\text{desired}}$  and step timing  $T$  have a slight inverse correlation, and so not all combinations of  $v_{\text{desired}}$  and  $T$  need to be generated. Most of the complexity of the action space lies in the choices for the parabola orientation,  $\phi$ . The choices for  $\phi$  are defined relative to the orientation of the velocity vector  $(\dot{x}, \dot{y})$  from the end of the previous footstep, and the discretization of  $\phi$  ranges from almost straight to almost U-turns. We note that the first choice that real humans would consider when navigating is to step directly toward the local goal. To address this, we create a special option for  $\phi$  that would orient the character directly toward its goal. With this specialized goal-dependent option, we found it was possible to give fewer fixed options for  $\phi$ , focusing on larger turns. Without this option, even with a large variety of choices for  $\phi$ , the character appears to steer toward an offset of the actual goal and then takes an unnatural corrective step.

### 4.2. Short-Horizon Best-First Search

We use a best-first search planner for a sequence of footsteps that minimizes energy cost. The implementation of our planner is the same as an  $A^*$  search, except for the *horizon*, described below.

The cost of taking a step is computed using Equations 2–5. The heuristic function used by the best-first search,  $h(s)$ , estimates the energy cost from the current state to a local goal:

$$h(s) = c_{\text{expected}} \times n \quad (6)$$

where  $c_{\text{expected}}$  is the energy spent in taking one normal footstep action based on the character's user-defined parameters, and  $n$  is the number of steps it would take to travel directly to the goal.

The *horizon* of our planner is the maximum number of nodes to be expanded for a single search. In most cases, a path is found before this threshold. We limit the horizon so that difficult or unsolvable situations will not cause a significant delay. If the planner searches too many nodes without reaching the goal, we instead construct a path to a node from the closed list that had the best heuristic value (the same closed list used in  $A^*$ ). Intuitively, this means that if no path is found to the goal within the search horizon, the planner returns a path to the reachable state that had the most promise of reaching the goal. The short-horizon approach guarantees that we will have at least some path for the character to use, even in difficult or unsolvable planning problems. In worst case, if no good solution is found, the path will simply be a sequence of 'stop' actions. For example, this can occur when a character is stuck dense environment. Eventually when the density clears, the character will continue.

### 4.3. Local Goals and Collision Avoidance

To navigate through large environments, we first plan a path using  $A^*$  (a traditional spatial path, not footsteps). Whenever a character needs to plan more footsteps, a local footstep goal is chosen, placed approximately 10 m ahead on the spatial path. This 10 m requirement is not strict; we experimented with other methods of choosing a local footstep goal, and they all worked decently well. Characters that are visible to each other can read each other's plans in order to predict their dynamic collision bounds at any given time. Visibility is determined by (1) having line-of-sight between the two characters, and (2) being within the character's visual field, modeled as a hemisphere centered around the character's forward-facing direction. This knowledge is analogous to the unspoken communication that occurs between real human pedestrians that makes human steering very robust. When a character re-plans, it does not try to avoid characters that it does not see, and therefore other characters, who are still executing old plans, may collide. The number of collisions can be drastically reduced by re-planning  $n$  steps in advance, before the previous plan is fully completed. This way there is always a 'buffer' of 2 or 3 steps that are guaranteed to be correct when a character predicts how to steer around another character. While deadlocks and collisions are still possible with this scheme, collisions are



**Figure 3.** (Left) A character side-steps and yields to the other pedestrian, then precisely navigates through the narrow doorway. (Right) An egress simulation. Characters do not get stuck around the corners of the glass door.

very rare, and we have not yet encountered a deadlock in our experiments.

#### 4.4. Choosing Exact Footstep Orientation

As described above, the planner maintains an interval of valid foot orientations for every step, constrained by the previous step's interval, as well as the trajectory of the current step. Once a sequence of footsteps has been planned, it is possible to choose exact footstep orientations. We constrain the interval of valid orientations once more using the *next* step's trajectory, now that this information is available. This computation relies on the same interval arithmetic described in Section 3. It is easy to see by contradiction that this process will not cause an invalid interval of orientations: if the interval becomes invalid during this postprocess, that would imply that no orientation of the current step could have produced a valid interval of the next step—but if this is true, that option would have already been pruned during planning and would not be encountered here. The exact orientation can be any value within this final interval; we found a good heuristic is to orient the foot as closely as possible to the orientation of the step's trajectory, with a special case for large turns.

## 5. RESULTS

For most results, characters are modeled with a center of mass 1 m above the ground, with a step length between 0.1 m and 1.0 m, step timing between 0.2 seconds and 0.8 seconds, and torso width of 60 cm.

Our short-horizon planner can solve challenging situations such as potential deadlocks in narrow spaces. Figure 3 depicts a challenging doorway situation. In many previous algorithms, characters would 'fight' at the doorway and may reach deadlock. In our method, the characters exhibit predictive cooperation, where one character steps aside. The doorway, 70 cm wide, is barely wide enough to fit a single pedestrian. In this tight situation, vector-based techniques would rely on collision prevention at the walls until the character eventually finds the door.

Our collision model allows tighter spacing in crowded conditions. An example is shown in Figure 3, where a group of characters squeeze through a glass door. With a single coarse collision radius, there would be many false-positive collisions. Instead, like real humans, these

characters are comfortable placing their feet and shoulders close to others in the dense crowd.

Our planner works online, in real-time. Performance is shown in Table 1, measured on a Core 2 processor, using a single thread. Planning is fast because of the scope of footsteps: a short horizon plan of 5–10 footsteps takes seconds to execute but only a few milliseconds to compute. The amortized cost of updating a character at 20 Hz is also shown in Table 1.

## 6. DISCUSSION AND FUTURE WORK

Footsteps are an appropriate form of control since they are the major contact point between a bipedal system and the external environment. By generating space–time sequences of footsteps, and by considering tighter dynamic collision bounds, our approach is able to control characters more precisely than existing crowd navigation techniques.

A 'stop' step is a specialized action in our planner. Being based on general planning, our technique can extend to use other specialized actions, such as running, jumping, even motion capture clips, as long as the action has well defined transitions, costs, and constraints. Existing steering techniques can also be emulated, for example, social forces models can be mapped to cost functions used by our planner.

There are some prominent aspects of bipedal locomotion which should be addressed in future work. Knee joints, ankle joints, muscles, angular momentum, and the center of pressure (pendulum pivot) shifting from heel-to-toe during a step—all of these affect the energy cost of real footsteps. We would also like to explore social and cognitive costs, where a character's objective may not necessarily be to minimize effort.

**Table 1.** Performance of our footstep planner for a character. The typical worst case plan generated up to 5000 nodes and expanded about 3000 nodes.

	Egress	2-way hall	700 boxes
	50 agents	200 agents	500 agents
Avg. # nodes generated	137	234	261
Avg. # nodes expanded	82	190	192
Planner performance	1.6 ms	4.4 ms	3 ms
Amortized cost 20 Hz	0.037 ms	0.1 ms	0.11 ms

## ACKNOWLEDGEMENTS

Authors thank Intel Corp. for their generous support through equipment and grants.

## REFERENCES

- van de Panne M. From footprints to animation. *Computer Graphics Forum* 1997; **16**(4): 211–223.
- Chung S-K, Hahn JK. Animation of human walking in virtual environments. In *Computer Animation*, 1999; 4–15.
- Coros S, Beaudoin P, Yin KK, van de Panne M. Synthesis of constrained walking skills. *ACM Transactions on Graphics* 2008; **27**(5): 1–9.
- Wu C-C, Medina J, Zordan VB. Simple steps for simply stepping. In *ISVC (1)*, 2008; 97–106.
- van Basten BJH, Peters PWAM, Egges A. *The Stepspace: Example-Based Footprint-Driven Motion Synthesis*. In *Computer Animation and Virtual Worlds, CASA 2010 Special Issue, Vol 21, Issue 3-4*, Chichester, UK: John Wiley and Sons Ltd., 2010.
- Autodesk. *3ds Max*, 2010.
- Helbing D, Molnár P. Social force model for pedestrian dynamics. *Physical Review E* 1995; **51**(5): 4282–4286.
- Reynolds CW. Steering Behaviors for Autonomous Characters, in the proceedings of Game Developers Conference 1999 held in San Jose, California. Miller Freeman Game Group, San Francisco, California. 1999. 763–782.
- Kapadia M, Singh S, Hewlett W, Faloutsos P. In Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, 2009. 215–223.
- Singh S, Kapadia M, Hewlett W, Reinman G, Faloutsos P. In Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, 2011.
- Pelechano N, Allbeck JM, Badler NI. Controlling individual agents in high-density crowd simulation. In *SCA, Eurographics Association*, 2007; 99–108.
- Gayle R, Sud A, Andersen E, Guy SJ, Lin MC, Manocha D. Interactive navigation of heterogeneous agents using adaptive roadmaps. *IEEE Transactions on Visualization and Computer Graphics* 2009; **15**(1): 34–48.
- Boulic R. Relaxed steering towards oriented region goals. In *MIG'08*, 2008; 176–187.
- van den Berg JP, Lin MC, Manocha D. Reciprocal velocity obstacles for real-time multi-agent navigation. In *ICRA*, 2008; 1928–1935.
- Lamarche F, Donikian S. Crowd of virtual humans: a new approach for real time navigation in complex and structured environments. *Computer Graphics Forum* 2004, **23**(3): 509–518.
- Paris S, Pettré J, Donikian S. Pedestrian reactive navigation for crowd simulation: a predictive approach. In *Eurographics 2007*, Vol. 26, 2007; 665–674. <http://www.morganclaypool.com/doi/abs/10.2200/S00123ED1V01Y200808CGR008>
- Feurtey F. Simulating the collision avoidance behavior of pedestrians. *Master's Thesis*, The University of Tokyo, School of Engineering, 2000.
- Badler N. *Virtual Crowds: Methods, Simulation, and Control*. Morgan and Claypool Publishers, 2008. DOI: 10.2200/S00123ED1V01Y200808CGR008 <http://www.morganclaypool.com/doi/abs/10.2200/S00123ED1V01Y200808CGR008>
- Paris S, Donikian S. Activity-driven populace: a cognitive approach to crowd simulation. *IEEE Computer Graphics and Applications* 2009; **29**(4): 34–43.
- Musse SR, Thalmann D. A model of human crowd behavior. In *Proceedings of the CAS'97*, Springer Verlag, Wien, 1997; 39–51.
- Shao W, Terzopoulos D. Autonomous pedestrians. In *SCA*, 2005; 19–28.
- van Basten BJH, Egges A. Path abstraction for combined navigation and animation. *MIG'09*, 5884/2009, 2009; 182–193.
- Lau M, Kuffner JJ. Precomputed search trees: planning for interactive goal-driven animation. In *SCA*, September 2006; 299–308.
- Kuffner, J.J., Jr., Nishiwaki K, Kagami S, Inaba M, Inoue H. Footstep planning among obstacles for biped robots. In *IEEE Intelligent Robots and Systems (IEEE/RSJ)*, Vol. 1, 2001; 500–505.
- Nishiwaki Kh, Sugihara T, Kagami S, Inaba My, Inoue S. Online mixture and connection of basic motions for humanoid walking control by footprint specification. In *ICRA*, Vol. 4, 2001; 4110–4115.
- Kuffner J, Nishiwaki K, Kagami S, Kuniyoshi Y, Inaba M, Inoue H. Online footstep planning for humanoid robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, September 2003.
- Li T-Y, Chen P-F, Huang P-Z. Motion planning for humanoid walking in a layered environment. In *ICRA*, Vol. 3, 2003; 3421–3427.
- Chestnutt J, Lau M, Cheung KM, Kuffner J, Hodgins JK, Kanade T. Footstep planning for the honda asimo humanoid. In *ICRA*, April 2005.
- Choi MG, Lee J, Shin SY. Planning biped locomotion using motion capture data and probabilistic roadmaps. *ACM Transactions on Graphics* 2003; **22**(2): 182–203.
- Zhang L, Pan J, Manocha D. Motion planning and synthesis of human-like characters in constrained environments. *MIG'09*, 5884/2009, 2009; 138–145.
- Wu J-C, Popović Z. Terrain-adaptive bipedal locomotion control. *ACM Transactions on Graphics* 2010; **29**(4): 72:1–72:10.
- Kuo AD. The six determinants of gait and the inverted pendulum analogy: a dynamic walking perspective. *Human Movement Science* 2007; **26**(4): 617–656.

## AUTHORS' BIOGRAPHIES



**Shawn Singh** is currently working on his Ph.D. at the University of California, Los Angeles. He received his M.S. in computer science from the University of Southern California. His research includes real-time photon mapping, novel forms of computation, and robust virtual pedestrian steering behaviors.



**Mubbasir Kapadia** received his B.E. in Computer Engineering in 2007 from University of Mumbai, India. He is currently working on his M.S. at the University of California, Los Angeles. His current research is applying egocentric approaches to pedestrian simulation and the evaluation of agent steering behaviors.



**Glenn Reinman** is an assistant professor in the department of computer science at UCLA. He received his B.S. from MIT in 1996 and his PhD and M.S. in Computer Science from UCSD in 2001. His main area of research is microprocessor architecture, and he directs the MARS lab at UCLA.



**Petros Faloutsos** is an assistant professor at the Department of Computer Science at the University of California at Los Angeles. He received his PhD degree (2002) and his MSc degree in Computer Science from the University of Toronto, Canada and his BEng degree in Electrical Engineering from the National Technical University of Athens, Greece. Professor Faloutsos is the founder and the director of the graphics lab at the Department of Computer Science at UCLA. The lab, called MAGIX (Modeling Animation and GrafIX), performs state of the art research in all aspects of graphics, focusing on virtual actors, virtual reality, physics-based animation and motor control. Professor Faloutsos is also interested in computer networks and he has co-authored a highly cited paper on the topology of the Internet. Professor Faloutsos is a member of the Editorial Board of the Journal Of The Visual Computer and has served as a Program Co-Chair for the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation. He is a member of the ACM and the Technical Chamber of Greece.