

Chapter 1

Virtual Tawaf: A Velocity-space-based Solution for Simulating Heterogeneous Behavior in Dense Crowds

Sean Curtis and Stephen J. Guy and Basim Zafar and Dinesh Manocha

Abstract We present a system to simulate the movement of individual agents in large-scale crowds performing the Tawaf. The Tawaf serves as a unique test case; the large crowd consists of a heterogeneous set of pilgrims, varying in both physical capacity and activity. Furthermore, the density of the crowd reaches extremely high levels (up to eight people/m².) This extreme density can place impractical constraints on simulation parameters. We use a velocity-space-based pedestrian model which exhibits consistent results even under extreme density: reciprocal velocity obstacles (RVO). Furthermore, we extend RVO to include *priority* and *right of way* — agents respond to potential collisions asymmetrically depending on context; one agent may yield, to varying degrees, to another. Our system uses a finite state machine to specify the behavior of the agents at each time step, to model the varied behaviors seen during the Tawaf. The finite-state machine, used in conjunction with RVO, generates collision-free trajectories for tens of thousands of agents in the performance of the Tawaf. The overall system can model agents with varying age, gender and behaviors, supporting the heterogeneity observed in the performance of the Tawaf, even at high densities.

Sean Curtis

University of North Carolina at Chapel Hill, Chapel Hill, NC, U.S.A., e-mail: seanc@cs.unc.edu

Stephen J. Guy

University of North Carolina at Chapel Hill, Chapel Hill, NC, U.S.A., e-mail: sjguy@cs.unc.edu

Basim Zafar

Hajj Research Institute, Umm al-Qura University, Makkah, Saudi Arabia, e-mail: Bzafar@uqu.edu.sa

Dinesh Manocha

University of North Carolina at Chapel Hill, Chapel Hill, NC, U.S.A., e-mail: dm@cs.unc.edu

1.1 Introduction

The Tawaf is one of the Islamic rituals of pilgrimage performed by Muslims when they visit Al-Masjid al Harām. Located in Makkah, Saudi Arabia, Al-Masjid al Harām surrounds the Kaaba, the site Muslims around the world turn towards while performing daily prayers. Al-Masjid al Harām is the largest mosque in the world and is regarded as Islam's holiest place. During the Tawaf, Muslim pilgrims circumambulate the Kaaba seven times in a counterclockwise direction, while in supplication to God.

The Tawaf is performed both during the Umrah and the Hajj. Performing the Hajj is one of the five pillars of Islam and every Muslim aspires to visit Makkah at least once in his or her life. Annually, more than two million Muslims perform the Hajj. While the Hajj has several stages and takes place over several days, all pilgrims move through the various stages of the Hajj on the same days which creates limitations in both time and space resulting in very high crowd densities during the Tawaf, especially on the Mataf, the marble floor of the mosque, in the center of which stands the Kaaba. During the Hajj season, or the last few days of the month of Ramadan, as many as 35,000 pilgrims perform Tawaf at the same time in the Mataf area in Al-Masjid al Harām. Given the large scale of the gathering, it is important to understand and model the behavior and movement of the crowd to provide insight which may improve crowd management techniques and help ensure the safety of the pilgrims.

The Tawaf has several properties which make simulating it particularly challenging:

Heterogeneous Population: The population of pilgrims varies significantly, spanning gender, a wide range of ages and physical capacity, and representing different cultures from all over the world.

High Density: The crowd density throughout the Mataf often varies considerably. It can become as high as eight pilgrims per square meter near the Kaaba [32]. The extremely high density greatly restricts the movement of the pilgrims.

Varying Velocities: The velocity of the pilgrims in Mataf can vary depending on many factors such as their distance from the Kaaba and the proximity of structures on the floor or congestion caused by other agents due to the capacity saturation and geometry of the mosque; the irregular shape of the Mataf is not well suited to the inherently elliptical movement around the Kaaba.

Complex Motion Flows: Different types of crowd flows have been observed during the Tawaf. These flows arise out of the sometimes contradictory intentions of the many pilgrims; at any given time, pilgrims will be simultaneously trying to stand still to kiss the Black Stone at the corner of the Kaaba, circumambulate the Kaaba, or attempt to move orthogonally to the circular flow, inwards, toward the Kaaba, or outwards, towards the exit, preventing purely circular flow.

Simulating the Tawaf will afford those who administer Al-Masjid al Harām the ability to evaluate alternative crowd flow control systems or architectural changes to improve the comfort and safety of the pilgrims and increase the capacity of the

Mataf. But creating a practical simulator for such a complex scenario is challenging. The crowd simulator must account for the heterogeneous population, allowing for large variance in the capabilities and actions of the agents. Furthermore, to capture the acts of the ritual, the simulator must provide a mechanism in which the activities and strategies of each agent change with respect to time. These features must be embedded in a computationally efficient simulator. It should scale well with respect to both the number of virtual pilgrims and increasing density. The greater the computation time, the less flexibility the simulator provides in evaluating scenario variations or producing stochastic studies in which multiple runs with randomly perturbed initial conditions are analyzed in aggregate. Satisfying these challenges and producing an accurate simulation of complex and dynamic interactions between pedestrians of this sort remains an open problem.

Main Results In this paper, we describe a system to model the movement of individual agents in a large-scale crowd performing the Tawaf. To address the above challenges, we present an agent-based model which combines a velocity-based pedestrian model to control local interactions between the agents with a finite state machine (FSM) to model the intentions of each pilgrim. We extend the pedestrian model with a parameter called *priority* which governs how the agents divide the effort to avoid collision between them. In some cases, agents act cooperatively to avoid collision. In other cases greater priority gives one agent *right of way* over another agent, causing the agent with less priority to yield. This enables us to model the asymmetric relationships observed between pilgrims in the Tawaf such as when pilgrims stop to kiss the Black Stone while others move around them. To model the changing goals of agents, each state in the FSM encodes a particular behavior which defines both strategy and tactics for navigating the shared space. The state provides a function which defines time-dependent values for a sub-space of the agent configuration space, including, but not limited to, such agent properties as preferred velocity and priority. The pedestrian model computes a collision-free trajectory by computing a new velocity based on the agent’s time-dependent state. We use several criteria to transition between the states based on spatial, agent-property, temporal and stochastic conditions.

The resulting system allows us to simulate crowds of heterogenous individuals, including variations in age and gender, performing the Tawaf. Each agent is associated with a unique instance of the pilgrim FSM. The FSM defines the general form of the Tawaf ritual, but each individual instance can allow for individual variance in the particular performance. For example, some pilgrims may possess a strong desire to approach the Kaaba, while others avoid the dense region near the Kaaba and maintain a greater distance.

From these simulations, we measure aggregate behavior such as density and velocity. We also measure Tawaf-specific metrics, such as the time to complete the Tawaf, and the overall throughput, in terms of the number of pilgrims that can complete the Tawaf per hour and show correlation with empirical observations.

Paper Organization: The rest of the paper is organized as follows. In Section 1.2, we survey related work on crowd simulation, behavior modeling, and simulation of the Tawaf. We discuss the full simulation pipeline in Section 1.3, paying

particular attention to the formulation of the high-level behavioral finite state machine. In Section 1.4 we present four models for pedestrian simulation and discuss their particular suitability for modeling the Tawaf. In addition, we discuss the details of *priority* and *right of way* and illustrate its effect on pedestrian relationships. Section 1.5 contains specific details on the actual performance of the Tawaf by living pilgrims and the mapping to a particular finite state machine. Finally, in Section 1.6 we provide the results of our system.

1.2 Related Work

In this section, we discuss related work in crowd simulation and behavior modeling for crowds. We also highlight some prior crowd simulation systems designed for simulating the Tawaf.

1.2.1 Crowd Simulation

There is extensive literature on crowd simulation and many techniques have been proposed.

Cellular automata (CA) are some of the oldest approaches for crowd simulation. In CA the workspace of agents is divided into discrete grid cells which can be occupied by zero or one agent. Agents then follow simple rules to move towards their goals through adjacent grid cells [26].

Continuum methods such as [27] and [20] treat the crowd as a whole and model the motion and interactions of agents based on equations that represent aggregate flow.

Agent-based approaches model each individual in the crowd and the interactions between them. Different techniques have been proposed to model these interactions. Reynolds [24] proposed Boids, which is a simple method based on rules for avoiding collisions while preserving flock cohesion. The rules are often implemented as forces. Other well known force-based methods including the social force model [12] (and its many variations), generalized centrifugal force model [6] and HiDAC [23]. These approaches use more complex forces between agents to model a larger domain of local interactions. Ondřej et al. [21] proposed a vision-based model in which agents respond to nearby obstacles based on the angle to the obstacle and the estimated “time to interaction”. Recently, velocity-space methods have been proposed to model human pedestrians. These geometric formulations are often based on velocity obstacles [9, 5, 11] and have been shown to exhibit many emergent crowd phenomena.

1.2.2 Behavior Modeling

Many researchers have proposed approaches to simulate various aspects of human and crowd behaviors. Funge et al. [10] proposed using a cognitive model to allow agents to plan and perform high-level tasks. Yu and Terzopoulos [31] introduced a decision network framework that is capable of simulating interactions between multiple agents. Ulincy and Thalmann [28] used a modular behavioral architecture to allow a mixture of automated and scripted behavior in multi-agent simulations. Durupinar et al. [8] modeled the effects personality factors have on local behavior. Yersin et al. [30] used spatial patches to direct motion and behavior of agents. Bandini et al. [4] applied a state machine to an underlying CA model to create scenarios with more complex behaviors. Yeh et al. [29] employed a physical collision avoidance mechanism to model abstract factors in pedestrian interactions such as aggression, priority and authority.

Data-driven approaches have also been used to capture crowd behaviors, often by training models of agent motion based on video data. Lee et al. [16] used data-driven methods to create group behavior such as queueing and clustering. Ju et al. [14] proposed a data-driven method which attempts to match the style of simulated crowds to those in a reference video. Patil et al. [22] proposed a method of directing crowd simulations with flow fields extracted from video or specified by a user. Video data has also been used to analyze and interpret real-world crowd behavior. Mehran et al. [17] proposed a method to detect abnormal crowd behavior from video using the social-force model. Johansson et al. [13] used video to study crowd behavior during portions of the Hajj.

1.2.3 Tawaf Simulation

There is some prior work on simulating crowd movement during the Tawaf and other Hajj related rituals. Algadhi and Mahmassani [3] simulated crowd flows in the Jamarat area of the Hajj using continuum models. Mulyana and Gunawam [19] performed agent based simulations of various rituals of the Hajj including a 500-agent simulation of the Tawaf. Zainuddin et al. [33] used the commercial software SimWalk to perform a social force-based simulation of up to 1,000 agents performing the Tawaf ritual. Sarmady et al. [25] performed a large crowd simulation of the Tawaf using CA techniques combined with a discrete-event simulator.

A few studies have also been performed on crowd flow in the Mataf area in the Al-Masjid al Harām. Al-Haboubi and Selim [2] proposed a potential spiral movement path to increase safety and throughput of pilgrims during the Tawaf. Koshak and Fouda [15] collected trajectories of actual pilgrims performing the Tawaf during the Hajj using GPS devices. The Crystals project currently studies how to incorporate cultural differences into simulations of Hajj pilgrims [1].

1.3 Modeling Crowd Behaviors

In this section, we give an overview of our method for modeling the crowd behaviors during the Tawaf. Human behavior arises from the confluence of many factors, including culture, psychology, environment and physiology. Generally, human behavior spans a wide range of activity. When discussing *crowd behavior* we limit our discussion to those human behaviors which affect how humans share space. For example, two people standing and discussing current events are functionally equivalent to those same people negotiating a business deal; the topic is unimportant, but the fact that they are stationary at a fixed distance away from each other is the behavioral detail which most influences crowd simulation. We characterize behaviors which affect the crowd with two concepts: where does an individual wish to be and how do they interact with those around them in reaching their goal? The first deals with the agent's intention — the general strategy, such as what path to take through an environment. The latter addresses the immediate tactics applied to execute the strategy under the dynamic constraints of a populated environment.

1.3.1 Agent-based Simulations


To simulate the Tawaf, with its heterogeneous population and widely varying activities, we need an approach which can accommodate a high-level of per-agent variability. To that end, we model the crowd with individual agents. Each agent is characterized by its *physical state* (position, velocity, size, etc.), its *behavior state* (preferred velocity, priority, its FSM, etc.), and its *property set* (a collection of associated data appropriate to the scenario.) For example, a simulated pilgrim's property set includes a counter indicating how many circles the pilgrim has already completed around the Kaaba. The counter doesn't directly affect the computation of an agent's preferred velocity or how it interacts with other agents, but it is used in the behavior mechanism to know when the Tawaf is complete.

We model the behaviors of agents by coupling together a high-level finite-state machine (FSM) with a pedestrian model. The FSM evaluates the agent's physical state and defines the agent's behavior state and, optionally, changes values in the agent's property set. The behavior state is used by the pedestrian model, in conjunction with an agent's physical state, to compute a new velocity and update the agent's physical state. Figure 1.1 illustrates the two components of our system and how they interact.

1.3.2 The Behavior Finite State Machine

A finite state machine (FSM) defines the behavior of an agent at every time step. Each state in the FSM defines the behavior state, and, optionally, the property set

Fig. 1.1 We simulate crowd behaviors appropriate to the Tawaf by coupling a high-level finite-state machine (FSM) with a low-level pedestrian model. The FSM computes behavior-state and property-set values for each agent. The pedestrian model, in turn, updates the agent's physical state.



for the agent. By providing unique definitions, each state can impart a distinct, observable behavior on the agent. We do not require any particular method for a state to use to compute the behavior state for an agent. The choice is arbitrary. All that matters is that the values for the agent's behavior state produce the desired behavior. For example, one element of the behavior state is the preferred velocity. It may be computed in any number of ways: by a simple rule, or as the result of a complex algorithm using techniques as varied as guidance fields or roadmaps. We give specific examples of this in Section 1.5.3

We have classified the FSM's transitions into four categories based on the types of conditions which cause the transition to become active: spatial, property, temporal, and stochastic. A *spatial transition* will cause the agent's current state to change when the agent's position achieves some pre-defined spatial configuration, such as entering an area, leaving an area, etc. For example, this transition will signal the start or end of a circumambulation. The *property transition* moves the FSM from the current state to a new state if some element of the agent's property set conforms to a particular condition. In the Tawaf, this transition causes an agent to exit when it has completed 7 circles. The *temporal transition* acts as a timer for the state. The transition is activated when the agent has been in the current state for some pre-defined amount of time. For example, some agents in the Tawaf will stop and pray for a few seconds when completing a circumambulation. Finally, the *stochastic transition* becomes active according to a user-defined probability distribution. In the Tawaf, we expect that only a fraction of the participants stop to pray. We use the stochastic transition to model this distribution. Finally, we prioritize the transitions such that if two transitions conditions are both true, the transition with the higher priority is taken.

1.4 Pedestrian Modeling

With a mechanism in place to alter agent behavior over time, we need to select a pedestrian model to execute the high-level strategy. In this section, we discuss various types of pedestrian simulation algorithms and their suitability for a scenario such as the Tawaf. Finally, we present the reciprocal velocity obstacle pedestrian

model and describe *priority* and *right of way* — an extension which increases the space of interactions between virtual agents.

1.4.1 Models of Pedestrian Simulation

There are numerous algorithms for simulating pedestrians. Each has its own unique set of advantages and disadvantages. It has yet to be shown that any single algorithm perfectly models pedestrian dynamics in arbitrary scenarios. For our purposes, we are most interested in those algorithms well suited to a specific scenario: the Tawaf. We are interested in simulators which will provide a mechanism to model the physical and behavioral heterogeneity observed in the Tawaf. Furthermore, to be useful, we require the simulator to be efficient (a hypothetical “perfect” algorithm which took hours to simulate seconds of data would be impractical.)

We divide pedestrian simulators into two categories: macroscopic and microscopic. Macroscopic approaches model a crowd of pedestrians as an aggregate phenomenon (e.g. [27, 20].) Microscopic simulators deal with individual agents, trusting that the aggregate behavior will naturally arise from basic principles (e.g. [24, 12, 6, 5, 9, 21].)

Macroscopic models typically assume a relatively high density; they operate on the principle that the choices of an individual pedestrian are strongly constrained by its local conditions. In dense crowds this is a reasonable assumption. As the crowd becomes sparser, the effect of one pedestrian on its distant neighbors is significantly reduced. In some cases, these approaches can be quite efficient. Narain et al. [20] were able to simulate 100K agents at 450 ms per simulation step. These approaches usually treat the crowd as a continuous, homogenous medium. Assuming continuity and homogeneity precludes the variation in physical attributes and behaviors observed in the performance of the Tawaf. For example, such approaches would be unable to create a dense simulation in which some agents remain stationary while other agents move next to them. For these reasons, we consider macroscopic simulation algorithms to be inappropriate for simulating the Tawaf.

Microscopic models provide greater potential to realize the kind of per-agent heterogeneity we require. Each agent is individually simulated and, as such, can be assigned arbitrary properties to model varying physical capacity. Furthermore, their behaviors can be individually specified; one agent’s motion is not explicitly constrained by its neighbors. It can try to pursue a goal that stands in direct opposition to its neighbors (its success is dependent on the pedestrian model.) We consider three major categories of microscopic models: cellular automata (CA), social forces (SF) and velocity obstacles (VO).

As previously indicated, CA approaches decompose the simulation domain into a uniform grid. Each agent occupies a single cell and a single cell can contain at most one agent. Probabilities are applied to neighboring cells based on a movement protocol and each agent’s position is updated according to the probability distribution and a set of rules for resolving conflict. CA approaches are typically simple to

implement. A CA approach has even been applied to simulating the Tawaf before [25]. However, the authors indicate that while CA can generate emergent phenomena (such as lane formation, etc.), the individual microscopic trajectories are “unrealistic” [25]. Furthermore, the grid decomposition imparts a homogeneity on the agents as well. Agents can only move an integer number of cells in a single time step. Finally, the authors indicate that CA has limitations with respect to density. The maximum possible density is simply a function of the size of the cells. This maximum is only theoretical. In practice, that density is impossible to achieve because if all cells are filled, agents cannot resolve conflicts and the simulation reaches a deadlock. We feel that, despite its simplicity, the cost of the spatial discretization leads to too many undesirable artifacts to simulate the Tawaf effectively.

SF-based and VO-based approaches both operate in continuous space, obviating the artifacts observed in CA. SF-based models treat pedestrians as mass-particles. Various forces applied to a particle draw the particle towards a goal position and prevent collisions with obstacles and other agents. The many variations of SF-based models generally differ in how the forces are formulated.

VO-based approaches consider the relative velocities and positions of agents to select a feasible velocity — a velocity which will remain collision free for a specified window of time. For each neighboring agent, it computes a set of infeasible velocities — velocities which will lead to collision within a specified time window. The selected feasible velocity is the velocity which lies outside the union of all infeasible velocity regions but which minimizes some cost function. There are multiple variations on VO-based algorithms, which may differ in how they model the space of inadmissible velocities, define the cost function, and how they solve the optimization problem.

Generally, both SF- and VO-based algorithms appear to be viable candidates for simulating the Tawaf to the level of fidelity we seek. More detailed investigation is required to differentiate their suitability. We provide summaries of a recent SF-based model [6] and a recent VO-based model [5]. For simplicity, we limit the summary to agent-agent interactions and refer the reader to the original papers for details on agent-obstacle interactions.

Generalized Centrifugal Force: The Generalized Centrifugal Force (GCF) model is a SF-based model which formulates inter-agent repulsive forces in terms of the agents’ positions and velocities.¹ The agent is modeled with the state vector: $[m \ \mathbf{p} \ \mathbf{v} \ \mathbf{v}^0]^T \in \mathbb{R}^7$, where $m \in \mathbb{R}^1$ is the agent’s mass, $\mathbf{p}, \mathbf{v}, \mathbf{v}^0 \in \mathbb{R}^2$ are the agent’s current position, current velocity, and preferred velocity, respectively. Preferred speed $v^0 = \|\mathbf{v}^0\|$ is simply the magnitude of preferred velocity.

At each time step, agent i ’s acceleration is computed as:

$$a_i = \frac{\mathbf{F}_i}{m_i} = \frac{\mathbf{F}_i^{drv} + \sum \mathbf{F}_{ij}^{rep}}{m_i}, \quad (1.1)$$

¹ The velocity term is the inspiration for the name. The original SF model considered only agent positions [12].

where \mathbf{F}_i^{drv} is a “driving” force and \mathbf{F}_{ij}^{rep} is the repulsive force applied to agent i by agent j . Given this acceleration, the agent’s velocity and position are updated by integrating with respect to time using an explicit integrator.²

The driving force is what causes the agent to move toward its goal. These systems assume that for each time step, a “preferred velocity” (\mathbf{v}^0) is computed. How this velocity is defined is arbitrary, but we assume it represents the velocity the agent would “prefer” to take in the absence of dynamic constraints. The driving force is defined such that it imparts an acceleration on the agent sufficient to reach its preferred velocity in τ seconds:

$$\mathbf{F}_i^{drv} = m_i \frac{\mathbf{v}_i^0 - \mathbf{v}_i}{\tau}. \quad (1.2)$$

The presence of other agents may prevent an agent from following its preferred velocity. This interference is modeled by repulsive forces. Each nearby agent j applies a repulsive force to the agent i of the form:

$$\mathbf{F}_{ij}^{rep} = -m_i k_{ij} \frac{(\eta v_i^0 + v_{ij})^2}{d_{ij}} \hat{\mathbf{e}}_{ij}, \quad (1.3)$$

$$v_{ij} = \max(0, (\mathbf{v}_i - \mathbf{v}_j) \cdot \hat{\mathbf{e}}_{ij}), \quad (1.4)$$

$$k_{ij} = \max\left(0, \frac{\mathbf{v}_i \cdot \hat{\mathbf{e}}_{ij}}{v_i}\right) \quad (1.5)$$

where $d_{ij} = \|\mathbf{p}_j - \mathbf{p}_i\|$ is the distance between agents i and j , $\hat{\mathbf{e}}_{ij} = \frac{\mathbf{p}_j - \mathbf{p}_i}{d_{ij}}$ is the normalized *direction* vector from agent i to agent j , v_{ij} is the amount of agent i ’s and j ’s relative velocity that lies in the direction of $\hat{\mathbf{e}}_{ij}$, clamped to the range $[0, \infty]$, η is a simulation variable used to tune the behavior of the simulation, and k_{ij} is a field-of-view weight — the strongest response is to agents in the direction of travel with decreasing weight as the angle increases to 90° on either side of that direction.

According to the authors, the formulation of the repulsive force has several desirable properties:

1. Repulsion is a local effect because the magnitude of the force is dependent on inverse distance. effect.
2. The v_{ij} term accounts for relative velocity so that a slow moving agent will not be affected by a fast moving agent in front of it.
3. The k_{ij} term gives the agent an active field of view. Agents will not be repulsed by agents behind them.

According to (1.3), the repulsive force between agents has infinite support; no matter how far the distance between two agents, some small contribution to one agent’s acceleration will be due to an unreasonably distant agent. Conversely, when the agents overlap, their distance converges to zero and the repulsive force can grow


² While the formula doesn’t preclude using an implicit integration scheme, the common practice has been to use a low-order explicit integrator such as forward Euler.

infinitely large. The authors combat both of these undesirable artifacts by approximating (1.3) with a spline which bounds the growth at small distances and limits the functions domain to a user-defined maximum distance.

Unfortunately, this formulation still exhibits some undesirable properties as well. The combination of how the forces are defined and the integration scheme can lead to very “jittery” agent behavior, especially under high densities; an agent’s trajectory may exhibit high-frequency oscillations because of numerical integration error which can only be addressed through taking extraordinarily small simulation time steps.

The full analysis of this behavior is beyond the scope of this work. However, we feel a brief intuitive discussion of the causes will illustrate why we deem a social-force-based model impractical for simulating the Tawaf. We leave the full, formal analysis for future work. We focus on two particular properties of the formulation as the cause of the undesirable oscillatory behavior: the explicit integration of a “stiff” physical system and the out-of-phase nature of driving and repulsive forces.

Fig. 1.2 The inverse distance function. At small distances, a small perturbation in distance leads to a large change in the function value (red circles.) The same sized perturbation at large distances leads to a correspondingly small change in function value (blue diamonds.)



The inter-agent repulsive force (1.3) is essentially a function of the inverse distance between the agents (see Fig. 1.2.) The function has relatively compact support; the force is greatest at near distances when collision is most imminent. As distance increases, the magnitude quickly decreases. When simulating a relatively sparse environment, where distance between agents is high, the magnitude of the repulsive force is quite small. A small perturbation in distances between agents produces forces with only slightly different magnitude. But when agents are close, small changes in distance lead to very large changes in forces—in other words, the slope of the force function is quite steep. This is a classic characteristic of a stiff system. When performing explicit integration, the common practice for SF-based pedestrian models, small time steps must be taken to prevent oscillatory behavior and unbounded error in the undamped system.

Furthermore, the driving force (1.2) and repulsive forces (1.3) are solved out of phase with each other. Imagine two agents moving towards each other at their preferred velocity. Because they are moving at their preferred velocity, their driving force is zero. They continue on their trajectories until they are close enough for

the repulsive forces to be non-zero. At that time step when the repulsive forces are first non-zero, the driving force is still zero. Thus, the repulsive forces are the sole influence on the agent and the current velocity is accelerated accordingly. At the next time step, the repulsive force will be significantly reduced (because the relative velocity has been reduced), but the driving force now increases due to the deviation between preferred and current velocities. This alternating dominance can eventually converge to a steady-state where they will be in balance, but it requires a small time step.

We seek to simulate the Tawaf during its peak performance, when tens of thousands of pilgrims pack into a small area reaching densities as high as eight people per square meter. There is an unavoidable computational cost in increasing the size of the simulation to a 35,000 agents. If we also had to significantly reduce the simulation time step to an extremely small time step, the simulation would no longer be tractable in reasonable time frames. For this reason, we consider SF-based models impractical for simulating the Tawaf.

Reciprocal Velocity Obstacle: The Reciprocal Velocity Obstacle (RVO) updates an agent’s velocity by performing geometric calculations in velocity space. Agents in RVO are modeled with a state vector similar to that of GCF: $[\mathbf{p} \ \mathbf{v} \ \mathbf{v}^0]^T \in \mathbb{R}^6$, where \mathbf{p} , \mathbf{v} , and \mathbf{v}^0 are defined as before.




Fig. 1.3 Velocity obstacles. (a) The velocity obstacle formed by agent j on i . (b) The truncated velocity obstacle for time window τ formed by j on i . (c) Agent i assumes the full burden to avoid collision, assuming j ’s velocity will be unchanged; VO_{ji} is displaced by j ’s velocity.

The *velocity obstacle* lies at the core of these approaches. As the name implies, it is an obstacle, but rather than lying in workspace or configuration space, it lies in velocity space. For agents i and j , agent j induces a velocity obstacle on i , VO_{ji} , and i induces a symmetric velocity obstacle on j , VO_{ij} . The velocity obstacle is a cone, originating at \mathbf{p}_i , which tightly bounds the Minkowski sum of agent i ’s geometry with j ’s. Fig. 1.3(a) illustrates VO_{ji} for two agents with circular geometry. If the relative velocity between agents i and j remains within this cone, there will be an inevitable collision. In practice, we are only concerned with collisions that can occur within the next τ seconds. Including this term truncates the cone (as illustrated in Fig. 1.3(b).)

This obstacle is the space of *relative* velocities that lead to collision. A single agent cannot exert unilateral control over the relative velocity. If agent i assumes that j will not change velocity, i must take full responsibility for avoiding the collision. This is accomplished by translating VO_{ji} by j ’s velocity, as shown in Fig. 1.3(c).


This is the original velocity obstacle formulation, in which each agent assumes that every other agent is a non-responsive, dynamic obstacle [9].

However, this model is a poor representation of people because agents do, in fact, react to each other. This leads to two significant issues. First, the velocity obstacle is only valid based on the assumption that the other agent maintains a constant velocity. It is possible for both agents to pick a velocity outside of their respective velocity obstacles but the resulting relative velocity may place them on a collision course. Secondly, both agents overreact to their neighbors (because they falsely assume that the other will make no effort to avoid collision.) This can easily lead to oscillatory motion as the agents overreact in successive steps.

Van den Berg et al. proposed an alternate formulation to VO which addresses these issues: Optimal Reciprocal Collision Avoidance (ORCA) [5]. The truncated cone, VO_{ij} , is replaced with a half plane, $ORCA_{ij}$. This solves the first issue by defining the half planes of $ORCA_{ij}$ and $ORCA_{ji}$ to contain “mutually reciprocal” sets of velocities. That means there is no pair of velocities, selected from each agents admissible region, which will lead to a collision within τ seconds. Furthermore, the planes are defined such that the amount of change to the relative velocity required to avoid collision is evenly distributed between the two agents, removing the danger of overcompensation and oscillation.

Finally, ORCA provides an additional advantage. When an agent has multiple neighbors, the inadmissible velocities is the union of all velocity obstacles. For truncated cones, computing this region, and finding the best admissible velocity outside is complex and expensive. With half planes, the admissible velocities form a convex polygon. For a convex cost function, the optimal velocity can be computed in $O(n)$ time, for n ORCA half planes.

Fig. 1.4 The formulation of the ORCA half-plane. The various components of the definition are illustrated. The minimum change in relative velocity, \mathbf{u} , the direction of minimum change, $\hat{\mathbf{n}}$, and the resultant ORCA velocity obstacles, $ORCA_{ji}$ and the symmetric half plane $ORCA_{ij}$.



The ORCA half-plane can be constructed geometrically in the following manner. Assume agents i and j adopt velocities \mathbf{v}_i^{opt} and \mathbf{v}_j^{opt} , respectively, and that these velocities place them on a collision course (i.e. $\mathbf{v}_i^{opt} - \mathbf{v}_j^{opt} \in VO_{ji}$.) Let \mathbf{u} be the vector from $\mathbf{v}_i^{opt} - \mathbf{v}_j^{opt}$ to the closest point on the boundary of the velocity obstacle (see Fig. 1.4.) More formally,

$$\mathbf{u} = (\operatorname{argmin}_{\mathbf{v} \in \delta VO_{ji}} \|\mathbf{v} - (\mathbf{v}_i^{opt} - \mathbf{v}_j^{opt})\|) - (\mathbf{v}_i^{opt} - \mathbf{v}_j^{opt}) \quad (1.6)$$

is the minimum change in relative velocity between i and j necessary to guarantee no collision within τ seconds. To model the reciprocity, half of the minimum change is applied to each agent. So, we can define the ORCA velocity obstacle induced by agent j on agent i as:

$$ORCA_{ji} = \{\mathbf{v} | (\mathbf{v} - (\mathbf{v}_i^{opt} + \frac{1}{2}\mathbf{u})) \cdot \hat{\mathbf{n}} < 0\}, \quad (1.7)$$

where $\hat{\mathbf{n}}$ is the normalized direction of \mathbf{u} . Proof of the guarantees can be seen in the original paper [5]. Henceforth, we will refer to this model as RVO.

RVO is less prone to jitter oscillations than GCF. Whereas GCF computes a new velocity by integrating a stiff physical system with explicit integration, RVO computes a new velocity directly in velocity space. Unlike GCF, for a given agent state, RVO will produce the same feasible velocity, regardless of time step size. Like GCF, agent position is still integrated using explicit integration and has limits on the size of the time step. Theoretically, the time step must be strictly larger than the time window τ . In practice, RVO has been shown to be stable for time steps as large as 0.2 s [7].

1.4.2 Priority and Right of Way

One of the appealing properties of RVO for pedestrian simulation is its reciprocal nature. The idea that moving pedestrians will each make an effort to avoid collisions with others is consistent with anecdotal evidence. However, the model's exactitude in defining the reciprocity to be precisely half implies a precision that does not exist in nature. While, generally, the equal division of effort is a reasonable model of the most generic behavior, there are scenarios in which effort is not shared equally and the model becomes highly dissatisfying.

On a subway platform, pedestrians enter the platform, find a location to await the train and then stop. In navigating the platform, moving pedestrians typically move around those already waiting. After they've stopped in their chosen position, those following behind, must move likewise around them. At that moment, the pedestrian shifts paradigms from an expectation of full responsibility for avoiding collisions, to the expectation that other moving pedestrians will assume the responsibility to avoid collision with them.

In a more subtle vein, even when all pedestrians are moving, the burden isn't necessarily shared equally. It may be that some pedestrians are more conservative and more willing to give way to others. Some pedestrians may seem more aggressive or determined. Subtle social and psychological clues affect how people react to each other and shifts the distribution of responsibility for avoiding collision.

The ability to model asymmetry plays an important role in the simulation of the Tawaf as well. There are several instances in which asymmetrical responses are vital to reproduce observed behaviors. When pilgrims queue up to kiss the Black Stone, their relationship with other pilgrims in the queue is different than with those still

circling; they should not yield position to those still circling, but must cooperate, to some degree, with those in the queue. When actually kissing the Black stone, a pilgrim must ignore all other pilgrims, holding their position in front of the stone. Finally, when exiting the Mataf, the agents must work outwards when the rest of the agents are moving tangentially or even spiraling inwards. At any given moment, the exiting pilgrims are in the minority. The ability of the minority to move counter to the flow of the majority is predicated on their ability to enforce their will on the majority.

In the study of traffic, there is a concept that perfectly captures this phenomenon: *right of way*. Right of way is the set of rules which define when one entity must yield to another entity. When moving pedestrians walk around standing commuters on a train platform, the stationary people have right of way. When an aggressive person moves through a crowd and those around him part to let him through, it is because he implicitly has right of way.

Unlike with vehicles, where right of way has a very discrete, exclusionary interpretation (i.e. between two cars, right of way belongs entirely to one vehicle), between pedestrians it can be considered a continuous quantity. Right of way can be absolute, when one pedestrian completely yields to another or it can be shared such that each pedestrian partially yields, albeit to different degrees, to avoid collision.

RVO's formulation provides a simple mechanism by which we can model continuous right of way. We introduce a new agent state parameter, p , called *priority*—a non-negative, real number. An agent with higher priority has right of way. We define the right of way of agent i over agent j as:

$$R_{ij} = \begin{cases} \max(1, p_i - p_j) & \text{if } p_i \geq p_j \\ 0 & \text{otherwise} \end{cases} . \quad (1.8)$$

As implied by (1.8), the value of R_{ij} lies in the range $[0, 1]$, regardless of what the relative priorities of the two agents are. Furthermore, $R_{ij} > 0$ implies $R_{ji} = 0$. Right of way can only be held by a single agent and an agent cannot have more than 100% right of way. Another implication of this formulation is that agents can be assigned tiered priorities — an aggressive agent may acquire full right of way over a passive agent, but it may still be required to yield right of way to a stationary agent. This is easily achieved by assigning priority values to the shy, aggressive and stationary agents of 0, 1, and 2, respectively (or any sequence of monotonically increasing values such that each value is at least one greater than the previous value.)

In the formulation of RVO, the velocity obstacle is defined with respect to an abstract relative velocity between agents i and j . The definition uses \mathbf{v}^{opt} to compute the relative velocity. Van den Berg et al. refer to this as the “optimization” velocity and suggest that this is typically the agent’s current velocity because it minimizes the amount of change to the current agent state required to avoid collision, but it need not necessarily be the current velocity [5].

We redefine \mathbf{v}^{opt} in terms of right of way. This new definition will affect the definition of ORCA $_{ji}$ (1.7) in the following manner:

$$ORCA_{ji} = \{\mathbf{v} | (\mathbf{v} - (\mathbf{v}_{ij}^{opt} + \alpha_{ij}\mathbf{u}_{ij}) \cdot \hat{\mathbf{n}} < 0)\}, \quad (1.9)$$

$$\mathbf{v}_{ij}^{opt} = (1 - R_{ij})\mathbf{v}_i + R_{ij}\mathbf{v}_i^0, \quad (1.10)$$

$$\alpha_{ij} = \begin{cases} 0.5 & \text{if } R_{ij} = R_{ji} = 0 \\ \frac{R_{ij}}{2} & \text{if } R_{ij} > 0 \\ \frac{1+R_{ji}}{2} & \text{if } R_{ji} > 0 \end{cases}. \quad (1.11)$$

The effect of right of way is as follows. If both agents have the same priority, no agent has right of way and the new formulation or RVO is equivalent to the old; both agents optimize with respect to their current velocities and share an equal burden in avoiding collision. As one agent's priority increases, its right of way also increases. The increased right of way affects the computation in two ways. First, the higher-priority agent's optimization velocity becomes a linear interpolation between its current and its preferred velocity. Second, the higher-priority agent's share of the burden linearly decreases. When an agent has full right of way, its optimization velocity is its preferred velocity and it bears no responsibility for avoiding collision.




Fig. 1.5 Four experiments for evaluating right of way. In each experiment, the white agent's progress is measured. (a) Experiment #1: A single agent moves through a stationary group of agents. (b) Experiment #2: A single agent holds position against a moving group of agents. (c) Experiment #3: A single agent moves perpendicularly to a moving group of agents. (d) Experiment #4: A circle of 100 agents, each trying to move to its antipodal position.

We illustrate the impact of priority and right of way in four experiments (see Fig. 1.5.) We apply the following methodology for each experiment. We construct a group of grey agents consisting of eight rows with 28 agents on each row. The rows are vertically offset to increase the average density. The priority of the grey agents always remains zero. We vary the priority of the white subject agent over the range $[0, 1]$. For each priority value, we run 20 iterations with a small random noise applied to the initial positions of the grey agents. In addition, for experiments 1, 2, and 3, we repeat the set of iterations while changing the average density of the grey agents over the values: 2, 3, 4, and 5 agents/m². Experiment 4 has a single density, 8 agents/m² (the maximum possible density when all agents converge in the center of the circle.) For experiments 1, 3, and 4, the subject agent travels from an initial position to a goal position. For these experiments, we measure the impact of priority by examining the travel time to its goal. More particularly, given its preferred speed (v^0) and the straight-line distance (d) to its goal, we compute the baseline travel time ($t_b = d/v^0$) and report the travel time as a multiple of the baseline. In experiment

2, the agent tries to maintain its position, so we examine the impact of priority by measuring the total distance it travels in the course of the simulation.³ The results of these experiments can be seen in Fig. 1.6 and Fig. 1.7.

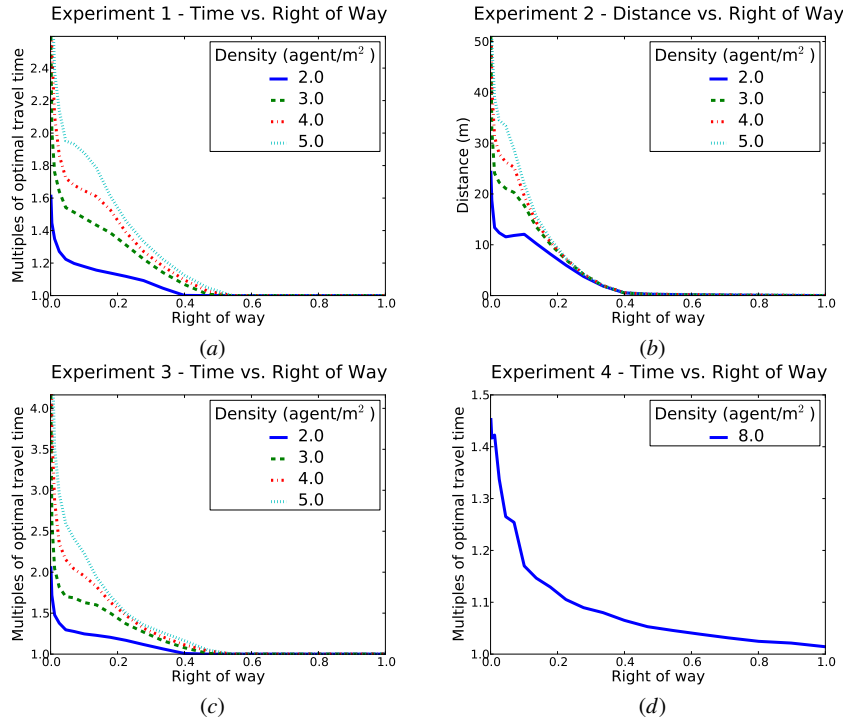


Fig. 1.6 The impact of priority on the experiment scenarios. (a), (c), and (d) report a multiple of the baseline travel time based on right-of-way value and density. (b) shows the absolute distance traveled.

There are several salient points to be made about the results of these experiments. First, in experiments 1, 2, and 3, as the subject agent's priority and the corresponding right of way increases, the subject agent's performance quickly converges to the baseline. This can be seen in Figs. 1.6(a), (b), and (c). The performance curves, at all densities, converge to the baseline value (bottom of the figure) at a priority value ranging between 0.4 and 0.6. This phenomenon becomes clearer when we observe the trajectory of the subject agent as shown in Fig. 1.7. The subject agent starts at the right in each figure and seeks to move in a straight line to its mirrored position on the left. The baseline trajectory would be a straight, horizontal line. With low priority, the agent is forced to deviate from the straight line. But for all priority levels, when the agent reaches the mid-point, it is able to travel directly toward its goal position.

³ If the agent were perfectly capable of maintaining its position, it would travel no distance at all.

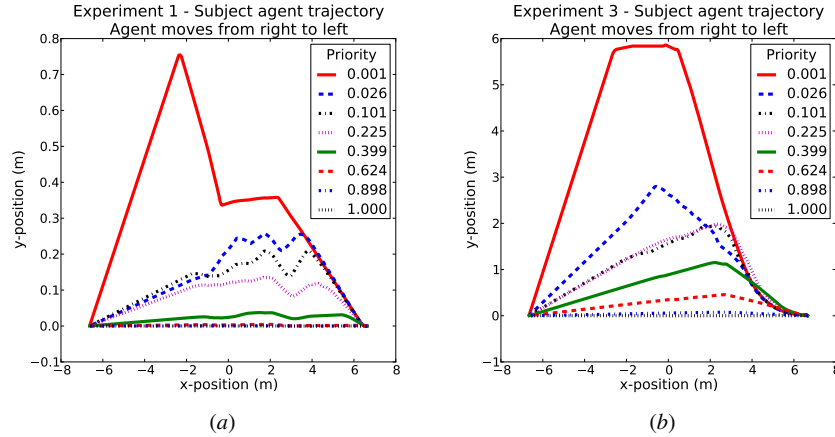


Fig. 1.7 The trajectory of the subject agent at varying priority levels. (a) Experiment 1. (b) Experiment 3.

We conjecture this quick convergence is due to two reasons. First, it has been shown that, like other pedestrian simulators, RVO exhibits emergent phenomena such as lane formation [11]. We conjecture that experiments 1, 2, and 3 benefit from this property. The experiments are orderly scenarios featuring simple bi-direction flows — an ideal circumstance for lane formation. The subject agent moves contrary to the large contingent of grey agents and as its priority increases, those agents nearest it begin to move out of its way. The following grey agents implicitly follow the divergent paths of the lead agents, forming lanes around the subject. Once those lanes have formed, the path for the subject agent remains clear. Second, the agents are arranged in a hexagonal lattice. Moving diagonally through the lattice is the clearest path possible. So, as the agent is pushed off of the horizontal, baseline trajectory, the most direct path to its goal eventually becomes a diagonal path which can exploit the greater clearance in the hexagonal lattice. So, for such orderly scenarios, a right-of-way value as little as 0.5 is sufficient for the subject agent to achieve baseline performance.

In comparison, experiment 4 represents a far more chaotic scenario. Agents moving to their antipodal positions do not share a preferred velocity with any of their neighbors. This significantly reduces the formation of lanes. The subject agent must contest with every agent in its path to achieve its goal. The experimental results support this idea. Fig. 1.6(d) shows increasing priority values contribute to the subject agent’s performance over the entire range of possible right-of-way values.

In addition, the impact of priority and right of way are dependent on the density around the subject agent. This is as expected. When the region around the subject agent is densely populated, taking any trajectory counter to its neighbors is significantly more difficult. The cause is two-fold. First, because the neighbors are near, the amount they interfere with the subject agent’s intentions is much higher; the subject is in danger of colliding with its neighbors in a very small time frame. Also,

nearby agents have very little flexibility in responding to the subject agent. So, the agent with right of way needs more priority to successfully influence its neighbors. But in a sparsely populated areas, neighboring agents are more distant, interfering less with the subject agent, and have a great deal more space to respond to the higher priority agent which leads to fast convergence to the baseline value. For the sake of visual clarity, we have vertically clipped the data shown in Fig. 1.6; the performance of the subject agent without right of way in high density scenarios was extraordinarily bad. Including those complete curves would have rendered the lower-density curves undifferentiable. At a density of $5 \text{ agents}/\text{m}^2$, the subject agent required 4.1X as much time for experiment 1, traveled 71.6 m in experiment 2, and took 7.9X as much time for experiment 3.


It is worth underscoring, that we are not modeling specific psychological factors nor advocating specific values which map human personality traits to priority values. That is a question for sociologists and psychologists to address. We simply provide a mathematical model which reproduces the phenomenon of asymmetric responses between pedestrians. Whence this asymmetry springs is an open question and we would hope that fellow scientists, better qualified to study these issues, will provide for us suitable characterizations for when such asymmetric responses occur and to what degree.

1.5 Simulating The Tawaf

In this section we give specific details on how the observed behaviors for performing the Tawaf are modeled.

Fig. 1.8 shows the layout of the Mataf area, the location where the Tawaf takes place including the Kaaba, Hateem and Maqam Ibrahim. The Hateem is a semi-circular structure which was originally part of the Kaaba when the Kaaba was rebuilt in A.D. 692. The Maqam Ibrahim is a structure of religious significance, to the northeast of the Kaaba.

Fig. 1.8 The layout of the Mataf area in the Al-Masjid al Harām. Pilgrims walk seven counter-clockwise circles around the Kaaba and Hateem. Each circle starts in front of the black stone (indicated as the start region.)



1.5.1 *The Rite*

The Tawaf is performed in the following manner:

1. Pilgrims enter the Mataf area and proceed towards the Black Stone. The Black Stone is located at the Kaaba's eastern corner. This landmark serves as the start and finish point of each circumambulation.
2. After reaching the region in front of the black stone, pilgrims perform Istilam, which can consist of kissing the Black Stone, touching the Black Stone with hands, or raising hands towards the Black Stone, all while saying *Tekbir*, "God is Great". On crowded days, only a small number of pilgrims will attempt to approach the Black Stone to kiss it. Those desirous to kiss the Black Stone will queue up near the southeast wall of the Kaaba. A pilgrim typically will only seek to kiss the Black Stone once, if at all.
3. The pilgrims walk, in a counter-clockwise direction, around the Kaaba and Ha-teem.
4. At the completion of each circumambulation, the pilgrims perform Istilam again.
5. At the end of the seventh circle, the pilgrims perform a short prayer outside the Mataf area, preferably in front of the Maqam Ibrahim or any convenient location in the mosque. A small number approach to kiss the Black Stone upon completion of the Tawaf.
6. Pilgrims exit the Mataf area. A recent study [32] has shown that 61% of the pilgrims exit the Mataf through the Safa exit in preparation for the next ritual.

1.5.2 *Population Characteristics*


One of the parameters of our simulation is the composition of the population. To that end we specify agent characteristics using population *classes*. Each population class defines a numerical distribution of values for a set of agent parameters. These values represent the physical capacity of the virtual pilgrims. The classes we use in simulating the Tawaf include the following parameters:

1. **preferred speed**: a normal distribution.
2. **maximum speed**: a normal distribution.

Properties not enumerated in a class (such as agent radius) are the same for all agents. We defined four agent classes to model both genders in two age categories ("old" and "young".) Agents are assigned a population class based on a user-defined distribution. The initial position of the agents is uniformly distributed in a circular area around the Kaaba. To achieve "steady-state" as quickly as possible, we set the agents randomly to have already completed some number of circumambulations (a uniformly distributed integer in the range [0, 7].) Finally, we force the flow into the Mataf to be equal to the flow out of the Tawaf by reintroducing each exiting agent into the system at a random entrance.

The space occupied by the human body can reasonably be bound with an ellipse with major and minor radii of 0.24 m and 0.15 m, respectively, with an area of 0.11 m². RVO uses circles to represent agents. A circle with a 0.19-meter radius has the same total area as the ellipse (as shown in Fig. 1.9.) We use this circle to model the pilgrims. Circles of this size can be optimally packed to yield a maximum density of 8 agents/m².

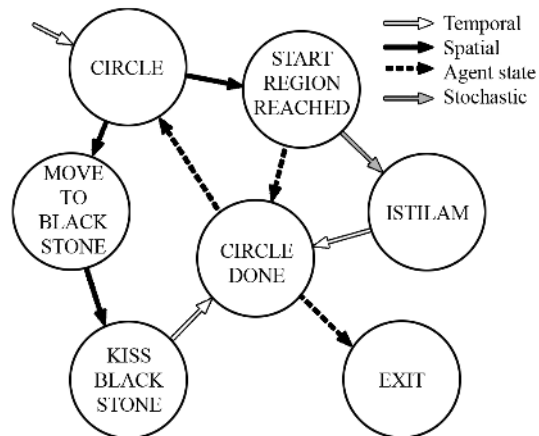
Fig. 1.9 A circle of radius 0.19 m has the same area as an ellipse with major and minor axes of 0.24 m and 0.15 m, respectively.



1.5.3 The Tawaf FSM

We have mapped the above behavior description to an FSM as shown in Fig. 1.10. Here we will enumerate the states and their transitions.

Fig. 1.10 The finite state machine for performing the Tawaf. Pilgrims start in the CIRCLE state. At the end of each circle, they either attempt to move to the black stone or perform Istilam and then perform another circle. After seven circles, they begin movement towards an exit.



CIRCLE: The circle state is the main circumambulation state. It contains two velocity components represented as guidance fields (a 2D vector field defined over the simulation domain specifying velocity directions.) The first is a radial guidance field with directions pointing towards the center of the Kaaba and the second is a tangential guidance field representing the direction of travel around the Kaaba. The

tangential field causes the pilgrims to circle around the Kaaba and the radial field draws them toward it. Although it is desirable to approach and kiss the Black Stone, on crowded days it can prove too difficult and many pilgrims choose not to attempt it. We model a variable degree of desire to approach the Kaaba and Black Stone by normally varying the weight of the radial velocity component. Agents with a large radial weight model those pilgrims with a greater desire to approach and put themselves in a better position to kiss the Black Stone.

There are two transitions out of this state. The first transition determines if an agent will queue up to kiss the Black Stone. The transition is a combination of spatial and property transitions. If the agent has not yet kissed the Black Stone and enters into a region near the southern corner of the Kaaba, the condition of the transition is met and the agent enters the MOVE TO BLACK STONE state.

The second transition is a spatial transition. If the agent reaches the start region in front of the Black Stone, the agent enters the START REGION REACHED state.

START REGION REACHED: This state is a decision point. It contains no velocity components. When an agent reaches this state, the state's transitions are evaluated and the agent immediately advances to the corresponding state.

This state contains two transitions. The first transition is a stochastic transition. This is the likelihood that a given agent will attempt to perform Istilam by stopping while turning to face the Kaaba. Anecdotal evidence suggests that this probability is about 15%. We generate a uniformly distributed random value in the range $[0, 1]$. If the value is in the range $[0, p]$, where p is the probability of stopping for Istilam, then the transition is active, moving the FSM to the ISTILAM state.

If the transition to ISTILAM is not taken, then the second transition is taken. This transition is, by definition, active. It moves the FSM to the CIRCLE DONE state.

CIRCLE DONE: This state is another decision point. Like START REGION REACHED, it contains no velocity components. At this state, we determine whether the agent has completed the Tawaf or not.

This state contains two transitions. The first transition is a property transition. If the agent has completed seven circles around the Kaaba, the FSM transitions to the EXIT state. Otherwise, the FSM transitions back to the circle state for the next circle.

MOVE TO BLACK STONE: This state controls the queue for those agents waiting to kiss the Black Stone. Upon entering this state, the agent is marked as having kissed the black stone. Subsequently, the transition from CIRCLE to MOVE TO BLACK STONE cannot be active for this agent. The velocity is computed as follows: the direction of the preferred velocity is towards the Black Stone. If there is another agent in the queue between the agent and the Black Stone, the speed is the lesser of two speeds: the agent's preferred speed and the speed that will guarantee the agent reaches the other agent's position in one second. If the space in front of the agent is clear, the preferred velocity's magnitude is simply the agent's preferred speed.

This state has a single spatial transition. It activates when the agent reaches the stone and moves the FSM to the KISS BLACK STONE state.

KISS BLACK STONE: This state contains a single velocity component and a single transition. Upon reaching the area directly in front of the Black Stone, the velocity is computed to hold the agent in that position. To aid in this purpose, the agent’s determination property is set to one. The single transition is a temporal transition. After a randomly determined duration the agent enters the CIRCLE DONE stage.

ISTILAM: This state, like the KISS BLACK STONE state, has a single velocity component and transition. It likewise computes a velocity to keep the agent fixed in the position at which the agent was when entering this state. However, this is a softer constraint and the determination is set to zero. The single transition is a temporal transition. After a randomly determined duration (1–2 seconds), the agent enters the CIRCLE DONE stage.

EXIT As pilgrims complete the Tawaf and exit the Mataf floor, they do so in a cooperative manner, continuing to circle the Kaaba and working their way towards the outside until they are in sufficient free space to head to their selected exit area. Each agent is randomly assigned an exit according to the probability distribution found in [32].

We have areas defined in the simulation domain for each of the five exits. Once the exit has been randomly selected, we then select a random point in the exit region to serve as the agent’s goal point.

To model the cooperative exit behavior exhibited by the pilgrims in the Tawaf, we generate the agent’s velocity with a weighted combination of three velocities: a vector from current position towards the exit goal position, a tangential component like that in the CIRCLE state, and an anti-parallel radial component (the opposite of the radial component of the CIRCLE state.) The tangential and anti-parallel radial components cause the agent to continue circling the Kaaba while working its way away from the Kaaba.

We blend the exit goal velocity and the circular velocity based on the agent’s local density. When the crowd is very dense, the agent continues around the Kaaba. As the local density reduces, the weight between goal and circular velocities changes linearly until an acceptable minimum density is achieved and the agent can move directly towards its end goal.

1.6 Results

We’ve run several simulations with our system. Our first goal is to achieve a result consistent with observed crowd movement during the Tawaf. To that end, we created a population of 35,000 agents with the following composition: 25% each of young male and female and 25% each of old male and female. Young males had a mean preferred speed of 1.0 m/s and a standard deviation of 0.2 m/s. Similarly old males had a mean preferred speed of 0.85 m/s with a standard deviation of 0.2 m/s. Young and old females had mean preferred speeds of 0.95 and 0.8 m/s, respectively. Both had a standard deviation of 0.15 m/s.

Our approach exploits the efficiency of the underlying pedestrian model. Our simulation used a time step of 0.1 s and was able to generate frames at 26 Hz on an Intel i7 running at 2.67 GHz. The evaluation of the FSM and pedestrian model were parallelized over the set of agents through the use of OpenMP. In essence, our simulator runs faster than real-time. For 35,000 agents, it produces 2.6 seconds of simulated results for each second of computation.

Fig. 1.11 The density of the crowd of pilgrims performing the Tawaf in our simulation. The dark region in the center is the Kaaba. Our simulation reaches a maximum density of 7.3 agents/m². The density field is computed as in [13].





Fig. 1.12 The speed of the individual agents performing the Tawaf in our simulation.



Figs. 1.11 and 1.12 show a single moment from our simulation results. In this image, approximately 25,900 agents are actively circling the Kaaba. The other 9,100 agents are entering, exiting or queuing to touch the Black Stone. The average walking speed of the circumambulating agents is approximately 0.73 m/s. The average completion time for the full Tawaf is 28.1 minutes. If we assume that the 25,900 circumambulating agents are representative of the portion of the population of 35,000 agents that are circling the Kaaba at any time, then this simulation implies a capacity of 55,300 participants per hour.

In 2008, Koshak and Fouda [15] tracked subjects performing the Tawaf with GPS devices. They partitioned the Mataf area into regions and computed the average speed for each region. The results of this analysis are shown in Fig. 1.13. We computed average speed for similar regions in our simulation. The simulated results can be seen in Fig. 1.14. The analysis shows that the simulation compares well with the real data in some ways and diverges in others.

1. Similarities

- a. Region 1, the region immediately preceding the start area, is the slowest region.
- b. Regions 5–7 exhibit higher speeds than regions 1–4.
- c. The top speed of the simulated crowd matches the top speed of the measured crowd.

2. Differences

- a. Simulated data exhibits a much narrower range of speeds.




Fig. 1.13 The observed speeds of real pilgrims traversing each region during the Tawaf [15].




Fig. 1.14 The average speed of simulated agents traversing each region during the Tawaf.

The disparity observed in the range of speeds can be attributed to two causes. First, when Koshak and Fouda performed their experiments, there was a line on the Mataf floor indicating the starting point. The line has since been removed. At the time, experts felt that as pilgrims approached the line, they would come to a stop while searching for the line. This is considered to be the dominant cause of the extreme slow down in the corresponding region. Our simulation models current behaviors reflecting the removal of the line. Thus, our agents don't come to a stop and the aggregate result is a higher speed through this region.

Secondly, the narrow range of simulated speeds may arise from properties in the pedestrian model. It may be that, as a pedestrian model, RVO is insensitive

to density-related effects, such as the fundamental diagram. Further research is required to confirm and correct, if necessary.

Heterogeneity: To explore the impact of the heterogeneous population, we ran two alternative simulations. One consisted of nothing but young males (the fastest pilgrim class.) The second simulation consisted solely of old females (the slowest pilgrim class.) The simulation consisting only of young males exhibited an average walking speed of 0.82 m/s for 24,900 circumambulating pilgrims and a corresponding Tawaf completion time of 25.5 minutes. In contrast, the simulation of old females obtained an average walking speed of 0.67 m/s for 26,500 circumambulating pilgrims with a Tawaf completion time of 30.2 minutes. The implied capacity is 58,600 pilgrims per hour for the young males and 52,700 pilgrims per hour for the old females.

The capacity indicated by the heterogeneous crowd is close to the average capacity of the two homogeneous crowds (although the heterogeneous crowd's capacity is slightly lower.) The full impact of heterogeneity is still unclear. It may be that populating the entire crowd with instances of a single, statistically average pedestrian may prove to be sufficient. This requires more study and requires better data concerning the demographics of the pilgrims performing the Tawaf and more flow data of the actual performance.

1.6.1 Limitations

While the results are promising, there are still aspects of the Tawaf it does not capture. In addition to the unknown impact of heterogeneity, these simulations haven't modeled groups. We currently treat the agents as individuals. To more fully capture the dynamics of the Tawaf, we would require a group model such as in [18]. There is, in particular, one instance of group behavior that has often been noted by observers. At times, a group of participants will force their way orthogonally across the crowd flow to get closer to the Kaaba. This behavior, its rate of incidence, and its characteristics are not well understood and, as such, is not included in our model.

More generally, the simulated speeds need to be validated. Although the maximum speed matches that observed by Koshak and Fouda [15], it may prove that at many of the densities observed, the speed of the pedestrians should be lower. Furthermore, since the time of Koshak's and Fouda's experiments, the Mataf area has been changed to improve the flow. We need to validate against more current data collected from coordinated GPS devices and cameras.

1.6.2 Conclusion

The unique nature of the Tawaf exhibits behaviors which are not well modeled by many existing crowd simulation systems. We have presented a framework for sim-

ulating many of the complex behaviors and relationships exhibited by pilgrims performing the Tawaf. By coupling a high-level finite-state machine with a low-level pedestrian model, we have been able to model a range of behaviors such as: circum-ambulating the Kaaba, queuing to touch the Black Stone, entering and exiting the Mataf floor, and pausing to perform Istilam. We've shown how to extend a velocity-obstacle-based pedestrian model to capture asymmetric inter-agent responses and have shown that the model is well behaved even at the extreme densities observed in our simulation. In many important respects, the results of the simulation match those observed in real people performing the Tawaf.

There are still multiple avenues to pursue for future work. The first is to confirm the validity of the pedestrian model with respect to density-dependent effects. In addition, we plan to extend the current set of behaviors to capture the important behaviors currently missing from our simulation, with particular focus on the impact of exiting pilgrims and groups. In addition, we intend to investigate the possibility of using video of the Tawaf to refine the behavior system, both the parameters and structure of the FSM as well as the local collision avoidance parameters.

Acknowledgements

This research is supported in part by ARO Contract W911NF-10-1-0506, NSF awards 0917040, 0904990 and 1000579.

References

1. Crystals project. <http://www.csai.disco.unimib.it/CSAI/CRYSTALS/>
2. Al-Haboubi, M., Selim, S.: A design to minimize congestion around the ka'aba. *Computers & industrial engineering* **32**(2), 419–428 (1997)
3. Algadhi, S., Mahmassani, H.: Modelling crowd behavior and movement: application to makkah pilgrimage. *Transportation and Traffic Theory* **1990**, 59–78 (1990)
4. Bandini, S., Federici, M., Manzoni, S., Vizzari, G.: Towards a methodology for situated cellular agent based crowd simulations. *Engineering societies in the agents world VI* pp. 203–220 (2006)
5. van den Berg, J., Guy, S.J., Lin, M., Manocha, D.: Reciprocal n-body collision avoidance. In: *Inter. Symp. on Robotics Research* (2009)
6. Chraïbi, M., Seyfried, A., Schadschneider, A.: Generalized centrifugal-force model for pedestrian dynamics. *Phys. Rev. E* **82**(4), 046,111 (2010)
7. Curtis, S., Snape, J., Manocha, D.: Way portals: Efficient multi-agent navigation with line-segment goals. *Proc. of Symposium on Interactive 3D Graphics and Games (I3D)* (2012)
8. Durupinar, F., Pelechano, N., Allbeck, J., Gudukbay, U., Badler, N.: The impact of the ocean personality model on the perception of crowds. *Computer Graphics and Applications, IEEE* **31**(99) (2010)
9. Fiorini, P., Shiller, Z.: Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research* **17**(7), 760–762 (1998)
10. Funge, J., Tu, X., Terzopoulos, D.: Cognitive modeling: knowledge, reasoning and planning for intelligent characters. In: *SIGGRAPH*, pp. 29–38. ACM Press (1999)

11. Guy, S.J., Chhugani, J., Curtis, S., Lin, M.C., Dubey, P., Manocha, D.: Pedestrians: A least-effort approach to crowd simulation. In: Symposium on Computer Animation. ACM (2010)
12. Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. *Physical Review E* **51**(5), 4282–4286 (1995)
13. Johansson, A., Helbing, D., Al-Abideen, H., Al-Bosta, S.: From crowd dynamics to crowd safety: A video-based analysis. *Advances in Complex Systems* (2008)
14. Ju, E., Choi, M.G., Park, M., Lee, J., Lee, K.H., Takahashi, S.: Morphable crowds. *ACM Trans. Graph.* **29**(6), 140 (2010)
15. Koshak, N., Fouda, A.: Analyzing pedestrian movement in mataf using gps and gis to support space redesign. In: The 9th International Conference on Design and Decision Support Systems in Architecture and Urban Planning (2008)
16. Lee, K.H., Choi, M.G., Hong, Q., Lee, J.: Group behavior from video: a data-driven approach to crowd simulation. In: Symposium on Computer Animation, pp. 109–118 (2007)
17. Mehran, R., Oyama, A., Shah, M.: Abnormal crowd behavior detection using social force model. *CVPR* (2009)
18. Moussad, M., Perozo, N., Garnier, S., Helbing, D., Theraulaz, G.: The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS ONE* **5**(4), e10,047 (2010). DOI 10.1371/journal.pone.0010047. URL <http://dx.doi.org/10.1371/journal.pone.0010047>
19. Mulyana, W., Gunawan, T.: Hajj crowd simulation based on intelligent agent. In: Computer and Communication Engineering (ICCE), 2010 International Conference on, pp. 1–4. IEEE (2010)
20. Narain, R., Golas, A., Curtis, S., Lin, M.C.: Aggregate dynamics for dense crowd simulation. *ACM Trans. Graph.* **28**, 122:1–122:8 (2009). DOI <http://doi.acm.org/10.1145/1618452.1618468>. URL <http://doi.acm.org/10.1145/1618452.1618468>
21. Ondřej, J., Pettré, J., Olivier, A.H., Donikian, S.: A synthetic-vision based steering approach for crowd simulation. In: Proc. SIGGRAPH, pp. 123:1–123:9 (2010)
22. Patil, S., van den Berg, J., Curtis, S., Lin, M., Manocha, D.: Directing crowd simulations using navigation fields. *IEEE TVCG* pp. 244–254 (2010)
23. Pelechano, N., Allbeck, J., Badler, N.: Controlling individual agents in high-density crowd simulation. In: SCA07 (2007)
24. Reynolds, C.: Flocks, herds and schools: A distributed behavioral model. In: SIGGRAPH (1987)
25. Sarmady, S., Haron, F., Talib, A.: A cellular automata model for circular movements of pedestrians during tawaf. *Simulation Modelling Practice and Theory* (2010)
26. Schadschneider, A.: Cellular automaton approach to pedestrian dynamics - theory. *Pedestrian and Evacuation Dynamics* (2001)
27. Treuille, A., Cooper, S., Popović, Z.: Continuum crowds. In: ACM SIGGRAPH 2006, pp. 1160–1168. ACM (2006)
28. Ulicny, B., Thalmann, D.: Towards interactive real-time crowd behavior simulation. In: Computer Graphics Forum, vol. 21, pp. 767–775. Wiley Online Library (2002)
29. Yeh, H., Curtis, S., Patil, S., van den Berg, J., Manocha, D., Lin, M.: Composite agents. *Proc. of SCA* pp. 39–47 (2008)
30. Yersin, B., Maim, J., Pettré, J., Thalmann, D.: Crowd patches: populating large-scale virtual environments for real-time applications. In: I3D09, pp. 207–214. ACM (2009)
31. Yu, Q., Terzopoulos, D.: A decision network framework for the behavioral animation of virtual humans. In: Symposium on Computer animation, pp. 119–128 (2007)
32. Zafar, B.: Analysis of the Mataf - Ramadan 1432 AH. Tech. rep., Hajj Research Institute, Umm al-Qura University, Saudi Arabia (2011)
33. Zainuddin, Z., Thinakaran, K., Abu-Sulyman, I.: Simulating the circumambulation of the ka'aba using simwalk. *European Journal of Scientific Research* **38**(3), 454–464 (2009)