

Virtual Tawaf: A Case Study in Simulating the Behavior of Dense, Heterogeneous Crowds

Sean Curtis

University of North Carolina at Chapel Hill

seanc@cs.unc.edu

Basim Zafar

Hajj Research Institute, Umm al-Qura University

Bjzafar@uqu.edu.sa

Stephen J. Guy

University of North Carolina at Chapel Hill

sjguy@cs.unc.edu

Dinesh Manocha

University of North Carolina at Chapel Hill

dm@cs.unc.edu

<http://gamma.cs.unc.edu/LARGE>

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 crowd consists of a heterogeneous set of pilgrims, varying with respect to physical capacity as well as activity. Furthermore, the density of the crowd reaches very high levels. Our approach uses a finite state machine to specify the behavior of the agents at each time step in conjunction with a geometric, agent-based algorithm to specify how an agent interacts with its local neighbors to generate collision-free trajectories. 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.

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, the Al-Masjid al Harām surrounds the Kaaba, the place Muslims around the world turn towards while performing daily prayers. The mosque is the largest 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 cre-

ates limitations in both time and space resulting in very high crowd densities during the Tawaf. During the Hajj, or the last few days of Ramadan, as many as 35,000 pilgrims perform Tawaf at the same time in the Mataf area in the 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 At any given moment, different pilgrims may move with different purposes, such as, entering or exiting the Mataf area, circling the Kaaba, or pausing to pray.

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 [28]. 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.

Complex Motion Flows Different types of crowd flows have been observed during the Tawaf. 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.

Crowd simulation methods that rely on an assumption of constant pedestrian parameters cannot capture these complex phenomena. The goals of pedestrians and their relationships with others can change arbitrarily with time. Producing an accurate simulation of complex and dynamic interactions between pedestrians 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 finite state machine (FSM) to model the intentions of each pilgrim and a geometric collision-avoidance algorithm to control local interactions between the agents. Each state in the FSM encodes a particular behavior, computing a unique “behavior vector” consisting of such per-agent properties as preferred velocity. The collision-avoidance algorithm uses this behavior vector in conjunction with the agent’s physical state to compute a collision-free trajectory. 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 heterogeneous individuals performing the Tawaf. We model age and gender variation by varying the agent’s preferred speed and maximum speed. Typically, young people move more quickly than old and men move more quickly than women. Each agent is associated with a unique instance of the pilgrim FSM. This means the immediate goal of any agent is independent of its neighbors. Finally, we vary the parameters which control how they pursue their goals. For example, some pilgrims 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 information, such as the time to complete the Tawaf, and the overall throughput of the area, in terms of the number of pilgrims that can complete the Tawaf per hour.

Paper Organization: The rest of the paper is organized as follows. In Section 2, we survey related work on crowd simulation, behavior modeling, and simulation of the Tawaf. We present an overview of our simulation pipeline in Section 3 and in Section 4 we describe the specific FSM which recreates Tawaf behavior. Finally we present our preliminary simulation results and compare them to real-world Tawaf crowd data in Section 5.

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.

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 [22].

Continuum methods such as [23] and [17] 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 [20] 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 [9] (and its many variations), generalized centrifugal force model [5] and HiDAC [19]. These approaches use more complex forces between agents to model a larger domain of local interactions. Recently, velocity-space methods have been proposed to model human pedestrians. These geometric formulations are often based on (Reciprocal) Velocity Obstacles (RVO) [25] and have been shown to exhibit many emergent crowd phenomena [8].

2.2. Behavior Modeling

Many researchers have proposed approaches to simulate various aspects of human and crowd behaviors. Funge et al. [7] proposed using a cognitive model to allow agents to plan and perform high-level tasks. Yu and Terzopoulos [27] introduced a decision network framework that is capable of simulating interactions between multiple agents. Ulincy and Thalmann [24] used a modular behavioral architecture to allow a mixture of automated and scripted behavior in multi-agent simulations. Durupinar et al. [6] modeled the effects personality factors have on local behavior. Yersin et al. [26] 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.

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. [13] used data-driven methods to create group behavior such as queueing and clustering. Ju et al. [11] proposed a data-driven method which attempts to match the style of simulated crowds to those in a reference video. Patil et al. [18] 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. [14] proposed a method to detect abnormal crowd be-

havior from video using the social-force model. Johansson et al. [10] used video to study crowd behavior during portions of the Hajj.

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 [16] performed agent based simulations of various rituals of the Hajj including a 500-agent simulation of the Tawaf. Zainuddin et al. [29] 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. [21] 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 [12] 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].

3. Modeling Crowd Behaviors

In this section, we give an overview of our approach that is used to model 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 the 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. 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? The first characterizes the paths an agent would take, free of constraints. The latter affects how that path is modified in a populated environment.

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 heterogeneity. To that end, we model the crowd with individual agents. Each agent is characterized by its *physical state* (the agent’s position, velocity, size, etc.), its *behavior state* (an FSM), and its *property set* (a collection of associated data appropriate to the scenario.) For example, a simulated pilgrim’s




Figure 1. We simulate crowd behaviors appropriate to the Tawaf by coupling a high-level finite-state machine (FSM) with a low-level local collision avoidance (LCA) algorithm. The FSM computes preferred velocity and LCA properties for each agent. The LCA algorithm, in turn, updates the agent’s physical state.

property set includes a counter indicating how many circles the pilgrim has already completed around the Kaaba. The counter doesn’t 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 low-level, local-collision avoidance (LCA) mechanism for performing local motion planning. The FSM evaluates the agent’s physical state and property set to define the parameters for the LCA algorithm, which, in turn, updates the agent’s physical state. Figure 1 illustrates the two components of our system and how they interact.

We considered several LCA algorithms used in crowd simulation. Although, CA has been used to simulate large crowds[21], the agents’ speeds are constrained by the spatial discretization, limiting the heterogeneity of the crowd. Furthermore, CA’s movement rules can easily lead to unnatural retrograde motion, e.g. agents may be more likely to move backwards than hold position. Continuum methods work well with large, locally coherent groups, but are not designed to model heterogeneous behaviors of numerous individual agents. Social-force models [9, 5, 19] and geometric approaches [25] allow for heterogeneity because they plan in continuous space. Both approaches compute trajectories with respect to a preferred velocity for each agent. To model the changes in interactions between agents, both approaches use a set of parameters which can be altered to change the influences agents have on each other. However, we find the force magnitudes in the social force methods to be more difficult to tune. Furthermore, the forces can easily lead to stiff physical systems which require small simulation time steps. The geometric approach seems to produce a well-behaved simulation even with a relatively large time step. While we believe our formulation should work for both types of methods, we selected “reciprocal velocity obstacles” (RVO) [25] because of these last two reasons.

3.2. Local Collision Avoidance

RVO [25] performs local collision avoidance by computing a space of collision-free velocities for an agent with respect to a number of neighboring agents. The i^{th} agent, A_i , computes a half plane for each of its neighboring agents, A_j , (and a symmetric half plane for agent A_j with respect to agent A_i .) The approach offers theoretical guarantees for optimality. The configuration of these half planes determine how the effort to avoid collision is shared between the agents. We have extended the publicly available RVO library to introduce a *determination* parameter for each agent which extends the space in how collision avoidance effort is apportioned between the agents. When two agents have equal determination, they share the responsibility for avoiding collisions equally. When one agent has a greater determination value, more burden is placed on the other agent, allowing the more “determined” agent greater latitude in pursuing its preferred velocity.

3.3. 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 preferred velocity, LCA parameters, and, optionally, the property set for the agent. By providing unique definitions, each state can impart a unique, observable behavior on the agent. The exact method a state uses for computing preferred velocity is arbitrary, so long as the resultant preferred velocity produces the desired behavior. The resultant velocity can be the product of a simple rule, or the result of a complex algorithm using techniques as varied as guidance fields or roadmaps.

We’ve classified the FSM’s transitions into four categories based on the types of conditions which cause the transition to be active: spatial, property, temporal, and stochastic. A *spatial transition* will cause the 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,




Figure 2. 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.)

the transition with the higher priority is taken.

4. Simulating The Tawaf

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

Figure 2 shows the layout of the Mataf area, where the Tawaf takes place including the Kaaba, Hateem and Maqam Ibrahim. The Hateem is a semi-circular section 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.

4.1. 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 (a short prayer said facing the Kaaba). A small number of pilgrims 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 Hateem.
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, 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.

- Pilgrims exit the Mataf area. A recent study [28] has shown that 61% of the pilgrims exit the Mataf through the Safa exit in preparation for the next ritual.

4.2. 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 unique values for a set of agent parameters. Particularly, we define a numerical distribution for each property in the class. The classes we use in simulating the Tawaf include the following properties:

- preferred speed:** a normal distribution.
- 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.

A human can be reasonably bound with an ellipse with a semiaxis radii of 0.24 m and 0.15 m, respectively, and 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. 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.

4.3. The Tawaf FSM

We have mapped the above behavior description to an FSM as shown in Figure 3. Here we will enumerate the states and their transitions.

CIRCLE: The circle state is the main circumambulation state. It contains two velocity components represented as guidance fields (a 2D, spatially-embedded field of velocities). 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




Figure 3. 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.

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 [28].

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




Figure 4. 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 [10].

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.

5. 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 LCA algorithm. Our simulations 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 FSM and LCA evaluation 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.

Figures 4 and 5 show a single moment from our simu-




Figure 5. The speed of the individual agents performing the Tawaf in our simulation.

lation 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 [12] tracked subjects performing the Tawaf with GPS devices. They partitioned the Mataf area into regions and computed the average speed for each region. We computed average speed for similar regions in our simulation. The results can be seen in Figure 6. Our results match Koshak and Fouda’s data in several respects:

1. Region 1 is the slowest region.
2. Regions 5–7 exhibit higher speeds than regions 1–4.
3. The top speeds of the simulated crowd matches the top speed of the measured crowd.

It is worth noting that, Koshak and Fouda’s data covers a wider range of speeds (from 0.27–0.73 m/s.) 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. 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.

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




Figure 6. The average speed of simulated agents traversing each region during the Tawaf.

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.

5.1. 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 [15]. 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 [12], it may prove that at the densities observed, the speed of the pedestrians should be lower.

5.2. 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 simulating many of the complex behaviors exhibited by pilgrims per-

forming the Tawaf. By coupling a high-level finite-state machine with a low-level local collision avoidance algorithm, we have been able to model a range of behaviors such as: circumambulating the Kaaba, queuing to touch the Black Stone, entering and exiting the Mataf floor, and pausing to perform Istilam. In many important respects, the results of the simulation match those observed in real people performing the Tawaf.

In the future, we plan to extend the current set of behaviors to capture the important behaviors currently missing from our simulation. 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. Website. <http://www.csai.disco.unimib.it/CSAI/CRYSTALS/>.
- [2] M. Al-Haboubi and S. Selim. A design to minimize congestion around the ka'aba. *Computers & industrial engineering*, 32(2):419–428, 1997.
- [3] S. Algadhi and H. Mahmassani. Modelling crowd behavior and movement: application to makkah pilgrimage. *Transportation and Traffic Theory*, 1990:59–78, 1990.
- [4] S. Bandini, M. Federici, S. Manzoni, and G. Vizzari. Towards a methodology for situated cellular agent based crowd simulations. *Engineering societies in the agents world VI*, pages 203–220, 2006.
- [5] M. Chraïbi, A. Seyfried, and A. Schadschneider. Generalized centrifugal-force model for pedestrian dynamics. *Phys. Rev. E*, 82(4):046111, 2010.
- [6] F. Durupinar, N. Pelechano, J. Allbeck, U. Gudukbay, and N. Badler. The impact of the ocean personality model on the perception of crowds. *Computer Graphics and Applications, IEEE*, 31(99), 2010.
- [7] J. Funge, X. Tu, and D. Terzopoulos. Cognitive modeling: knowledge, reasoning and planning for intelligent characters. In *SIGGRAPH*, pages 29–38. ACM Press, 1999.
- [8] S. J. Guy, J. Chhugani, S. Curtis, M. C. Lin, P. Dubey, and D. Manocha. Pedestrians: A least-effort approach to crowd simulation. In *Symposium on Computer Animation*. ACM, 2010.
- [9] D. Helbing and P. Molnar. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282–4286, May 1995.
- [10] A. Johansson, D. Helbing, H. Al-Abideen, and S. Al-Bosta. From crowd dynamics to crowd safety: A video-based analysis. *Advances in Complex Systems*, 2008.
- [11] E. Ju, M. G. Choi, M. Park, J. Lee, K. H. Lee, and S. Takahashi. Morphable crowds. *ACM Trans. Graph.*, 29(6):140, 2010.
- [12] N. Koshak and A. Fouda. 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.
- [13] K. H. Lee, M. G. Choi, Q. Hong, and J. Lee. Group behavior from video: a data-driven approach to crowd simulation. In *Symposium on Computer Animation*, pages 109–118, 2007.
- [14] R. Mehran, A. Oyama, and M. Shah. Abnormal crowd behavior detection using social force model. *CVPR*, 2009.
- [15] M. Moussad, N. Perozo, S. Garnier, D. Helbing, and G. Theraulaz. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS ONE*, 5(4):e10047, 04 2010.
- [16] W. Mulyana and T. Gunawan. Hajj crowd simulation based on intelligent agent. In *Computer and Communication Engineering (ICCCE), 2010 International Conference on*, pages 1–4. IEEE, 2010.
- [17] R. Narain, A. Golas, S. Curtis, and M. C. Lin. Aggregate dynamics for dense crowd simulation. *ACM Trans. Graph.*, 28:122:1–122:8, December 2009.
- [18] S. Patil, J. van den Berg, S. Curtis, M. Lin, and D. Manocha. Directing crowd simulations using navigation fields. *IEEE TVCG*, pages 244–254, 2010.
- [19] N. Pelechano, J. Allbeck, and N. Badler. Controlling individual agents in high-density crowd simulation. In *SCA07*, 2007.
- [20] C. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH*, 1987.
- [21] S. Sarmady, F. Haron, and A. Talib. A cellular automata model for circular movements of pedestrians during tawaf. *Simulation Modelling Practice and Theory*, 2010.
- [22] A. Schadschneider. Cellular automaton approach to pedestrian dynamics - theory. *Pedestrian and Evacuation Dynamics*, 2001.
- [23] A. Treuille, S. Cooper, and Z. Popović. Continuum crowds. In *ACM SIGGRAPH 2006*, pages 1160–1168. ACM, 2006.
- [24] B. Ulicny and D. Thalmann. Towards interactive real-time crowd behavior simulation. In *Computer Graphics Forum*, volume 21, pages 767–775. Wiley Online Library, 2002.
- [25] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha. Reciprocal n-body collision avoidance. In *Inter. Symp. on Robotics Research*, 2009.
- [26] B. Yersin, J. Maim, J. Pettré, and D. Thalmann. Crowd patches: populating large-scale virtual environments for real-time applications. In *I3D09*, pages 207–214. ACM, 2009.
- [27] Q. Yu and D. Terzopoulos. A decision network framework for the behavioral animation of virtual humans. In *Symposium on Computer animation*, pages 119–128, 2007.
- [28] B. Zafar. Analysis of the Mataf - Ramadan 1432 AH. Technical report, Hajj Research Institute, Umm al-Qura University, Saudi Arabia, 2011.
- [29] Z. Zainuddin, K. Thinakaran, and I. Abu-Sulyman. Simulating the circumambulation of the ka'aba using simwalk. *European Journal of Scientific Research*, 38(3):454–464, 2009.