Right of Way:

Asymmetric Agent Interactions in Crowds

Sean Curtis[1] Basim Zafar[2] Adnan Gutub[3]

Dinesh Manocha[4]

[1,4]University of North Carolina at Chapel Hill

[2]Hajj Research Institute, Umm al-Qura University

[3]HajjCoRE, Umm al-Qura University

[1]CB #3174 Sitterson Hall

[1]UNC Chapel Hill, NC 27599-3175

[1]Tel. (+00)1 919 962 1870 Fax. (+00)1 919 962 1799

[1]email: seanc@cs.unc.edu

Abstract

Pedestrian models are typically represent interactions between agents in a symmetric fashion. In general, these symmetric relationship are valid for a large space of scenarios in which we simulate crowds. However, there are many cases in which symmetric response between agents are inappropriate, leading to unrealistic behavior or undesirable simulation artifacts. We present a novel formulation, called *right of way*, which provides a well-disciplined mechanism for modeling asymmetric relationships between pedestrians. We illustrate its impact by applying it to two different types of pedestrian models and showing its effect in multiple scenarios. Particularly, we show how it enables simulation of the complex relationships exhibited by pilgrims performing the Tawaf.

Keywords: pedestrian dynamics, crowd simulation, multi-agent models, asymmetric responses

Introduction

Models of pedestrian motion and interaction tend to be based on an aggregate, statistical mean behavior. Each simulated pedestrian perceives and responds to its environment and those response tend to be symmetrical; agents respond to each other in a uniform manner. As with all such models, this mean behavior approach is acceptable for a wide variety of scenarios. However, assumptions of symmetry can be incompatible with many meaningful scenarios we may wish to simulate. Furthermore, symmetry can lead to undesirable simulation artifacts. To account for these issues, we need a well-formulated and well-disciplined mechanism to introduce asymmetric responses.

It is important to emphasize that we are not discussing the subtly nuances which naturally arise from human variation. In reality, when two pedestrians share the space, their responses will not be perfectly symmetric; one pedestrian may be more agile, more tentative, more aggressive, etc. This is encompassed by the mean behavior abstraction of the model. We are focused on scenarios in which the asymmetry in responses is not a subtle nuance, but a dominant factor.

A typical subway station offers an excellent context in which to examine examples of such scenarios. Picture an empty concourse. Groups of pedestrians enter the concourse from varying directions, heading toward subway platforms. As they perceive each other, they adapt their planned paths of travel to avoid conflict. We track one group onto their destination subway platform. The platform already contains a number of passengers, standing and waiting for the train. The pedestrians entering the platform wend their way between the standing passengers. Each moving pedestrian finds a place to wait and comes to a stop. The remaining moving pedestrians, in turn, find paths around these newly waiting passengers. When the train arrives, the waiting passengers move towards the doors. As the doors open, the waiting pedestrians make space for the disembarking passengers. While this happens, a tardy commuter quickly enters the platform and moves to one of the waiting groups and proceeds to aggressively push through the waiting crowd. Finally, the waiting passengers enter the train and the train departs.

In the common scenario described above, multiple behaviors and inter-agent relationships are exhibited.

Symmetric response The agents moving through the empty concourse exhibit the basic symmetric relationships; each pedestrian reasonably assumes that other pedestrians will make an effort to avoid conflict and plans accordingly. This behavior is also exhibited by the waiting passengers as they move towards the door.

Adaptation to non-responsive agents The moving pedestrians recognize that the passengers already waiting on the platform will, in most cases, not move to accommodate them. It becomes the responsibility of the moving pedestrians to avoid the unresponsive stationary passengers and the responsive moving passengers while seeking their own goal. When a moving passenger reaches they goal, they change from a responsive moving pedestrian, to an unresponsive stationary passenger.

Social priority Through cultural and social convention, it is the practice to allow those disembarking the train to take precedence over those boarding. The waiting passengers make way for those exiting the train although, if space is available, passengers may be able to board and disembark simultaneously.

Aggression Finally, the tardy passenger exhibits aggression. An overtly aggressive pedestrian can "force" its way through a crowd, even without making physical contact. Other pedestrians can recognize aggressive properties and choose to yield ground to them in order to avoid conflict.

In the four examples above, only the first behavior can be modeled in a mean behavior model which assumes symmetric responses to potential conflict. To model these more sophisticated scenarios, we need to introduce a mechanism which will allow such asymmetric relationships to be realized.

We introduce a simple model for capturing asymmetric relationships. It is based on the traffic concept of "right of way". Our model has several desirable mathematical properties which allow us to capture all of the behaviors listed above (as well as additional scenarios) in a well-disciplined manner. Furthermore, we can use the right-of-way model to improve simulation results in several situations (such as flow through narrow passages.) The concept of right of way is sufficiently general that it can be applied to a number of pedestrian models to good effect; we show it applied to common force-based and velocity-based techniques. Finally, we show it's efficacy in a particularly challenging scenario: the Tawaf.

The Tawaf is one of the Islamic rituals of pilgrimage performed by Muslims when they visit Al-Masjid al Harām during the Umrah and the Hajj. 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, the marble floor of the mosque. The density of the pilgrims reaches as high as eight people/m^2 [1]. And yet, even in these densely packed scenarios, pilgrims are able to pursue contrary goals, producing discontinuities in the flow. Such discontinuities would be impossible using purely symmetric responses.

Paper Organization: In the remainder of this paper, we discuss related work in pedestrian simulation and pedestrian bheaviors, describe the right of way model and how it is incorporated into two different pedestrian models, analyze the impact of the right of way model in experimental benchmarks and practical scenarios and, finally, discuss its limitations and summarize its effect.

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.

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 [2].

Continuum methods such as [3] and [4] 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 [5] 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 [6] (and its many variations), generalized centrifugal force model [7] and HiDAC [8]. These approaches use more complex forces between agents to model a larger domain of local interactions. Ondřej et al. [9] 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 [10, 11, 12] and have been shown to exhibit many emergent crowd phenomena.

Behavior Modeling

There is a great deal of research in behavior modeling. Typically, these efforts have focused on higher-order behaviors, such as determining where an agent wants to go, under what circumstances its goal might change, etc. Such works include cognitive modeling [13], decision networks [14], scripted behaviors through modular behavior architecture [15], and personality factors [16]. These approaches are complimentary to the concept of right of way. They determine what an agent wishes to accomplish, and right of way influences how they interact with nearby agents in achieving that goal.

Right of way operates at a lower level of behavior. There have been work in capturing these kinds of behavior as well. One of the most common approaches seek to account for asymmetric responses based on direction (i.e. strong responses to agents in front, weak responses to those behind) as in [17, 7]. In addition, data-driven approaches have sought to learn inter-pedestrian behaviors from video such as in [18, 19, 20].

Priority and Right of Way

In the subway example scenario described in the introduction, the asymmetric behaviors illustrated all have a common trait. In each case, the behaviors manifested consisted of one agent giving way for another agent. The moving pedestrians gave way to the waiting pedestrians, the boarding passengers gave way to the disembarking passengers, and the average pedestrians gave way to the aggressive pedestrian.

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. The standing passengers have right of way over the moving pedestrians. The disembarking passengers have right of way over the boarding passengers. Even the aggressive agent has right of way, implicitly granted by the other pedestrians who seek to preemptively avoid conflict.

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.

We model right of way for pedestrians by introducing a new agent parameter: *priority*. Right of way of one agent over another is defined by their relative priority. Specifically, priority (p) is a non-negative, real-valued number. We define the right of way of agent i over agent j as:

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

This formulation has several properties. First, as implied by (1), the value of R_{ij} lies in the range [0, 1], regardless of what the relative priorities of the two agents are; an agent cannot have more than 100% right of way. Second, $R_{ij} > 0$ implies $R_{ji} = 0$; right of way can only be held by a single agent. Third, 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 average, 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.)

Applying Right of Way

Right of way plays two roles in a pedestrian model. The first role is simple and intuitive; right of way affects the distribution of effort between two agents. For example, in the perfectly symmetric case, the effort to avoid collision is evenly distributed between the two agents. As right of way of increases for agent i, its portion of the effort decreases and the burden of agent j increases.

The second role is more subtle. It addresses the issue of what the effort achieves. If we only considered the first role, an agent right of way would not have to exert effort to avoid collisions, but it would not guarantee that the agent is actually able to pursue its goal. An agent with 100% right of way should be better able to achieve its intention — a waiting passenger should be able to maintain its position and the aggressive agent should be able to cut through the crowd in its preferred direction of travel. In the symmetric behavior, the collision avoidance effort only considers current physical state. As agent i's right of way increases, the collision that agents i and j seek to avoid becomes less dependent on agent i's current state and more on agent i's preferred velocity. Simply put, if agent i has 100% right of way, it should be able to directly pursue its preferred velocity and other agents should actively move out of its way.

We will show how both roles of right of way can be incorporated into two common pedestrian paradigms: social forces and velocity obstacles. In both cases, agents are represented as two-dimensional disks with a common basic state vector: $[r\ \mathbf{p}\ \mathbf{v}\ \mathbf{v}^0]^T \in \mathbb{R}^7$, where $r \in \mathbb{R}^1$ is the agent's radius, $\mathbf{p}, \mathbf{v}, \mathbf{v}^0 \in \mathbb{R}^2$ are the agent's current position, current velocity, and preferred velocity, respectively. By convention, current speed v and preferred speed v^0 are simply the magnitudes of the corresponding vectors. When we refer to a particular agent i's properties, we will apply a subscript to the property (e.g. r_i, $\mathbf{v}_i^0$, etc.) Finally, we provide the following symbols for convenience: d_{ij} is the distance between the centers of agent i and j, $\hat{\mathbf{d}}_{ij}$ is the unit vector pointing from agent i to j, and r_{ij} is the combined radii of agents i and j. Additional, model-specific parameters will be introduced as needed.

Social Forces

While there are many variants of social forces, we select a straightforward variant for this discussion [21]. In this variant, the interaction between two agents is controlled by isotropic repulsion forces. Specifically, agent i experiences a repulsive force implied by nearby agent j formulated as follows:

$$\mathbf{F}_{ij} = A_i e^{((r_{ij} - d_i j)/B_i)} \hat{\mathbf{d}}_{ji}, \tag{2}$$

where A_i and B_i are simulation constants controlling the magnitude and fall-off of the repulsive force, respectively. It is common practice (and advocated by the authors [21]) to define the parameters A_i, B_i, and agent mass, m_i, globally for all agents (A, B, and m). In this case, the forces applied to agents i and j, and the resultant accelerations, are equal in magnitude and opposite in direction (i.e. $\mathbf{F}_{ij} = -\mathbf{F}_{ji}$ and $\mathbf{a}_{ij} = -\mathbf{a}_{ji}$).

Rather than thinking of the two accelerations independently, we can think about them as a single, relative acceleration: $\bar{\mathbf{a}}_{ij}$. This is the relative acceleration which must be imparted to the two-agent sub-system to prevent collisions. How the acceleration is distributed between the two agents is arbitrary. Right of way's first role will determine the distribution of the acceleration by redefining the parameter A in terms of right of way.

$$A_ij(R_{ij}, R_{ji}) = \begin{cases} A(1 - R_{ij}) & \text{if } R_{ij} > 0 \\ A(1 + R_{ji}) & \text{if } R_{ji} > 0 \\ A & \text{otherwise} \end{cases} \tag{3}$$

Right of way's second role is to increasingly insure that the agent with right of way is able to pursue its preferred velocity. We do this by redefining the force *direction*. By default, the direction of $\mathbf{F}_{ij}$ points away from agent j. If the force were oriented perpendicularly to agent j's preferred velocity, then 100% of the force would be applied to moving agent i out of agent j's preferred path. So, we redefine the force direction by interpolating between the original force direction and this perpendicular force direction, based on the right of way.

$$\bar{\mathbf{d}}_{ji}(R_{ij}, R_{ji}) = \begin{cases} slerp(\hat{\mathbf{d}}_{ji}, R_{\perp}\hat{\mathbf{v}}_i^0, R_{ij}) & \text{if } R_{ij} > 0 \\ slerp(\hat{\mathbf{d}}_{ji}, R_{\perp}\hat{\mathbf{v}}_j^0, R_{ji}) & \text{if } R_{ji} > 0 \\ \hat{\mathbf{d}}_{ji} & \text{otherwise} \end{cases}, \tag{4}$$

where $\hat{\mathbf{v}}_i^0$ is the direction of agent i's preferred velocity, $slerp$ is the spherical linear interpolation function (with R_{ij} serving as the blend value), and $R_{\perp}$ is a rotation matrix which produces an orthogonal vector such that the angle between that vector and $\hat{\mathbf{d}}_{ji}$ is less than 90°. In the special case where the agent with right of way wishes to hold still — a $\mathbf{0}$ preferred velocity — the rotated vector lies within 90° of the agent's current velocity. This causes the agent to pass on the nearest side of the stationary agent.

Velocity Obstacles

As with social forces, there are a number of approaches which use velocity obstacles for local planning. We will provide the details for applying right of way to one particular variant: Optimal Reciprocal Collision Avoidance (ORCA) [11]. Like the social-forces approach, the ORCA algorithm takes an agent's preferred velocity and local simulation state to compute a feasible, collision-free velocity which still allows the agent to make progress toward its goal. However ORCA uses a significantly different mechanism. Rather than applying forces, ORCA performs geometric optimization in velocity space. In velocity space, each nearby agent and obstacle applies a constraint on the velocities the agent can safely take. All of these constraints are considered simultaneously, and ORCA selects the feasible velocity closest to the agent's preferred velocity. A full discussion of how ORCA works is beyond the scope of this paper and we refer the reader to the original for specific details. However, we will briefly summarize a few principles to give context to how right of way is incorporated with the underlying model.

At each time step, agent i seeks to travel its preferred velocity. Each nearby entity (agents and static obstacles) limit the agent's ability to accomplish this. The effect of the nearby entity is modeled by the velocity obstacle it projects. ORCA defines this velocity obstacle as a half-plane. The half-plane represents a space of velocities which, if taken, will lead agent i into collision with the corresponding entity. For a static obstacle, computation of the velocity obstacle is simple because the future state of the static obstacle can be perfectly predicted. For dynamic entities, such as other agents, the prediction is a conservative estimate. Agent i assumes that a neighbor will take a velocity inside a space of velocities and avoids the full space. Meanwhile, the neighbor agent performs a symmetrical calculation with respect to agent i. Right of way will alter how we define these symmetrical spaces.

The space of velocities is computed as follows. Given an *abstract* relative velocity between agents i and j, ORCA determines if a collision is possible within some user-defined window of time (τ). If so, ORCA computes the smallest change to relative velocity which will prevent this collision. Finally, the required relative velocity change is evenly apportioned between agents i and j. The velocity spaces are half planes positioned and oriented so that as long as each agent selects a velocity outside the half plane, the resultant *relative* velocity will include the minimum change computed before.

Right of way's first role determines how the required change to the relative velocity is apportioned. If we assume that $\mathbf{u}_{ij}$ is the required change, then, in the symmetric case, agent i would have to account for $0.5\mathbf{u}_{ij}$ and agent j would account for the symmetric value $0.5\mathbf{u}_{ji}$. We use right of way to change the weight so as an agent's right of way increases, its portion of $\mathbf{u}_{ij}$ decreases. So, the agent's share becomes $\alpha_{ij}\mathbf{u}_{ij}$ where:

$$\alpha_{ij} = \begin{cases} \frac{1-R_{ij}}{2} & \text{if } R_{ij} > 0 \\ \frac{1+R_{ji}}{2} & \text{if } R_{ji} > 0 \\ 0.5 & \text{otherwise} \end{cases} . \tag{5}$$

Right of way's second role manifest's itself in the definition of the abstract relative velocity mentioned earlier. By default, the relative velocity that ORCA uses is the literal interpretation: $\mathbf{v}_i - \mathbf{v}_j$. The authors of ORCA state that this isn't not the only viable choice [11]. By using the current velocities, $\mathbf{u}_{ij}$ represents the smallest change to the agent's current physical state (i.e. minimum acceleration.) If we were to replace $\mathbf{v}_i$ with $\mathbf{v}_i^0$ then we would compute the minimum change required for agent i to travel its preferred velocity. This is precisely the goal of right of way. So, we define the relative velocity between agents i and j, $\mathbf{v}_{ij}$ as:

$$\mathbf{v}_{ij} = (1 - R_{ij})\mathbf{v}_i + R_{ij}\mathbf{v}_i^0. \tag{6}$$

The net effect of (5) and (6) is that if both agents have the same priority, no agent has right of way and the default symmetric behavior is in effect; both agents optimize with respect to their current velocities and share an equal burden in avoiding collision. As agent i's right of way increases, agent j's burden to avoid collision increases and the perceived collision is in the direction of agent i's preferred velocity.

Analysis and Results

In this section we will illustrate the impact of right of way through four abstract experiments and one complex scenario (the Tawaf.)

Experiments

We have designed four simple experiments (see Figure 1). We apply the following methodology to each experiment. We construct a group of grey agents consisting of eight ranks with 28 agents on each ranks. The ranks 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/m^2. Experiment 4 has a single density, 8 agents/m^2 (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 right of way 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 right of way by measuring the total distance it travels in the course of the simulation.[1] The results of these experiments can be seen in Figure 2, Figure 3, and Figure 4.

There are several salient points to be made about the results of these experiments. We'll address each simulator in turn. **ORCA**: 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 Figures 2(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

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

becomes clearer when we observe the trajectory of the subject agent as shown in Figure 4(a) and (b). 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.

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 [12]. 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. Figure 2(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 Figure 2; 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 agents/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.

Social Forces: I must add a similar analysis to the social forces right of way as soon as the plots are done. Assume it'll be shorter than the ORCA analysis because much will be overlapped and I'll make reference to previous points. There will, most likely, be some indication of the differences which arise from predictions. The forces are highly localized which means the agents only yield right of way at the last minute which leads to issues. Furthermore, in the densest scenarios, applying the full force causes significant instability and chaos (as can be seen when priority = 1. I'll probably go back and change the formula so that I don't apply $2A$ to the force but some fraction of that and then justify it here. Also mention that with this social forces model, the ability of one agent to move contrary to its neighbors is significantly constrained. It required a large priority value before the agent was able to work its way through a crowd.

Narrow Passages

Predictive models, such as ORCA, can have problems at narrow passages such as doorways. As one agent begins to move into the doorway, another agent, coming from behind and to the side of the first, may cause the first agent to back up to avoid a perceived collision. In normal circumstances, this wouldn't happen because the pedestrians can recognize that the person who has entered the doorway first will pass the portal before the other approaches. These simple models have no such knowledge. We can impart this knowledge by use of priority and right of way.

We create a simple navigation mechanism to facilitate flow through the doorway. We define a region immediately surrounding a doorway. While an agent is in this region, their priority grows at a rate commensurate with their distance to the doorway.

This helps implicitly prioritize the agents in the doorways. Those agents that have arrived near the doorway earlier than others will have a head start in accumulating priority and will be better able to pass through the portal, undisturbed by other agents.

We show some analysis of some flow through a narrow passage. We show that ORCA (and possibly SF) flow too slowly through the doorway (as compared to real human data). With the inclusion of this simple mechanism, agents are able to pass through the doorway more efficiently.

The Tawaf

Some more details on the Tawaf. What it is, the performance of the Tawaf. The specific details of the Tawaf which require right of way: queuing for the black stone and, possibly, exiting.

Make mention of continuum models here. They are predicated on the assumption that crowds of people, particularly *dense* crowds of people behave much like fluids. In dense crowds, individual choices matter less because the motion of others constrains individuals to move like their neighbors. This thesis has led to approaches such as [3] and [4]. As with most models, this abstraction captures some nice macroscopic behaviors but it can be shown that it is not universally true. In fact, in the performance of the Tawaf, it is precisely the ability of pilgrims to resist the flow which makes queuing possible. It is the exceptions that matter the most.

Summary and Limitations

Summary

Many practical and important real-world scenarios require pedestrian models which can capture asymmetric relationships between agents. We have provided a simple mathematical formulation for modeling such asymmetric relationships which we call "right of way". We have successfully applied right of way to two significantly, representative pedestrian models. In both cases, we were able to extend the space of interactions the pedestrian model was able to produce. Finally, we have shown how right of way has allowed us to simulate critical aspects of the Tawaf; pilgrims in incredibly dense circumstances are able to maintain a queue and effect an exit despite the significant flow.

Limitation

Mapping priority and right of way into the model can be difficult. As has been shown here, the incorporation of right of way into the social force model appears significantly different than in the velocity obstacle model. It may well prove that there is a better mapping of right of way to forces than has been provided here.

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.

References

[1] Basim Zafar. Analysis of the Mataf - Ramadan 1432 AH. Technical report, Hajj Research Institute, Umm al-Qura University, Saudi Arabia, 2011.

[2] Andreas Schadschneider. Cellular automaton approach to pedestrian dynamics - theory. *Pedestrian and Evacuation Dynamics*, 2001.

[3] A. Treuille, S. Cooper, and Z. Popović. Continuum crowds. In *ACM SIGGRAPH 2006*, pages 1160–1168. ACM, 2006.

[4] Rahul Narain, Abhinav Golas, Sean Curtis, and Ming C. Lin. Aggregate dynamics for dense crowd simulation. *ACM Trans. Graph.*, 28:122:1–122:8, December 2009.

[5] C.W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH*, 1987.

[6] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282–4286, May 1995.

[7] Mohcine Chraibi, Armin Seyfried, and Andreas Schadschneider. Generalized centrifugal-force model for pedestrian dynamics. *Phys. Rev. E*, 82(4):046111, 2010.

[8] N. Pelechano, JM Allbeck, and NI Badler. Controlling individual agents in high-density crowd simulation. In *SCA07*, 2007.

[9] Jan Ondřej, Julien Pettré, Anne-Hélène Olivier, and Stéphane Donikian. A synthetic-vision based steering approach for crowd simulation. In *Proc. SIGGRAPH*, pages 123:1–123:9, 2010.

[10] Paolo Fiorini and Zvi Shiller. Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research*, 17(7):760–762, July 1998.

[11] J. van den Berg, Stephen J. Guy, Ming Lin, and D. Manocha. Reciprocal n-body collision avoidance. In *Inter. Symp. on Robotics Research*, 2009.

[12] Stephen J. Guy, Jatin Chhugani, Sean Curtis, Ming C. Lin, Pradeep Dubey, and Dinesh Manocha. Pledestrians: A least-effort approach to crowd simulation. In *Symposium on Computer Animation*. ACM, 2010.

[13] J. Funge, X. Tu, and D. Terzopoulos. Cognitive modeling: knowledge, reasoning and planning for intelligent characters. In *SIGGRAPH*, pages 29–38. ACM Press, 1999.

[14] 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.

[15] B. Ulicny and D. Thalmann. Towards interactive real-time crowd behavior simulation. *Computer Graphics Forum*, 21(4):767–775, 2002.

[16] 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.

[17] A. Johansson, Helbing D., and P. K. Shukla. Specification of the social force pedestrian model by evolutionary adjustment to video tracking data. *Advances in Complex Systems*, 10:271–288, 2007.

[18] Kang Hoon Lee, Myung Geol Choi, Qyoun Hong, and Jehee Lee. Group behavior from video: a data-driven approach to crowd simulation. In *Symposium on Computer Animation*, pages 109–118, 2007.

[19] Eunjung Ju, Myung Geol Choi, Minji Park, Jehee Lee, Kang Hoon Lee, and Shigeo Takahashi. Morphable crowds. *ACM Trans. Graph.*, 29(6):140, 2010.

[20] S. Patil, J. van den Berg, S. Curtis, M.C. Lin, and D. Manocha. Directing crowd simulations using navigation fields. *IEEE TVCG*, pages 244–254, 2010.

[21] D. Helbing, I. Farkas, and T. Vicsek. Simulating dynamical features of escape panic. *Nature*, 407, 2000.

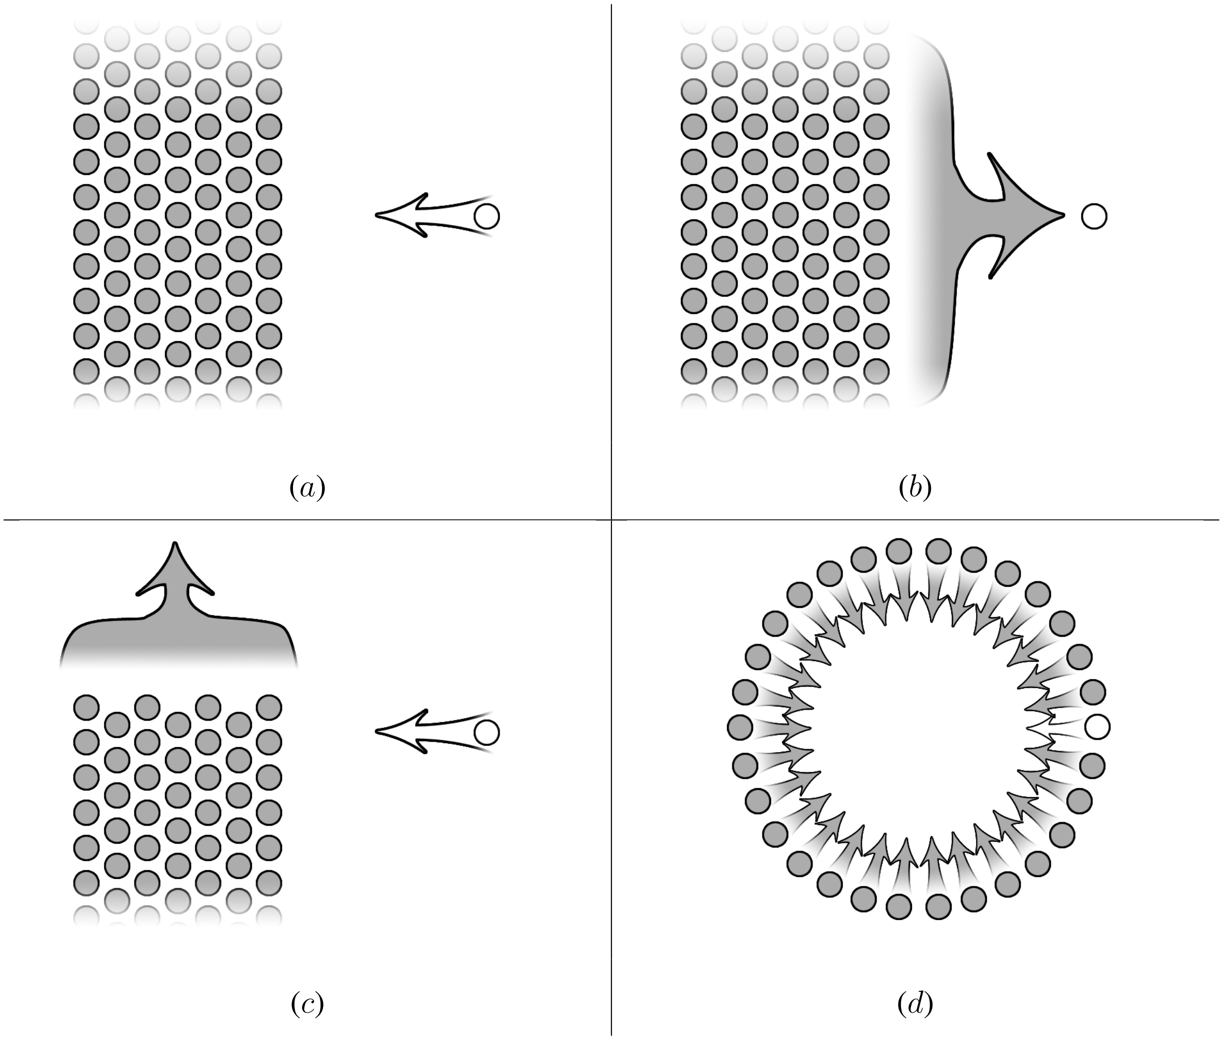

(a) (b) (c) (d)

Figure 1: 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.

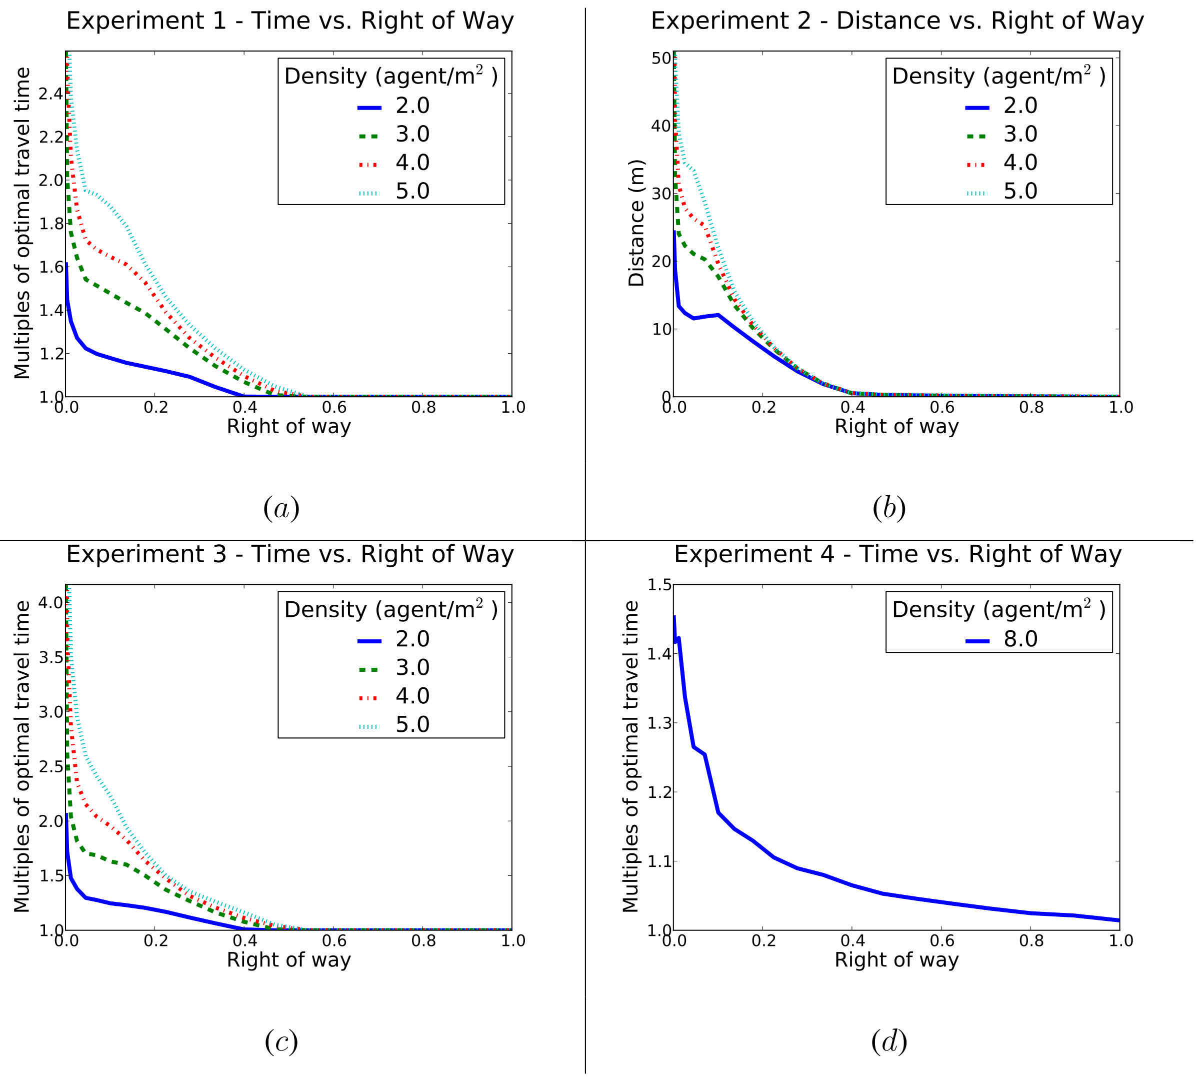

Figure 2: The impact of priority on the experiment scenarios with ORCA. (a), (c), and (d) report a multiple of the baseline travel time based on right-of-way value and density. (b) shows the absolute distance traveled.

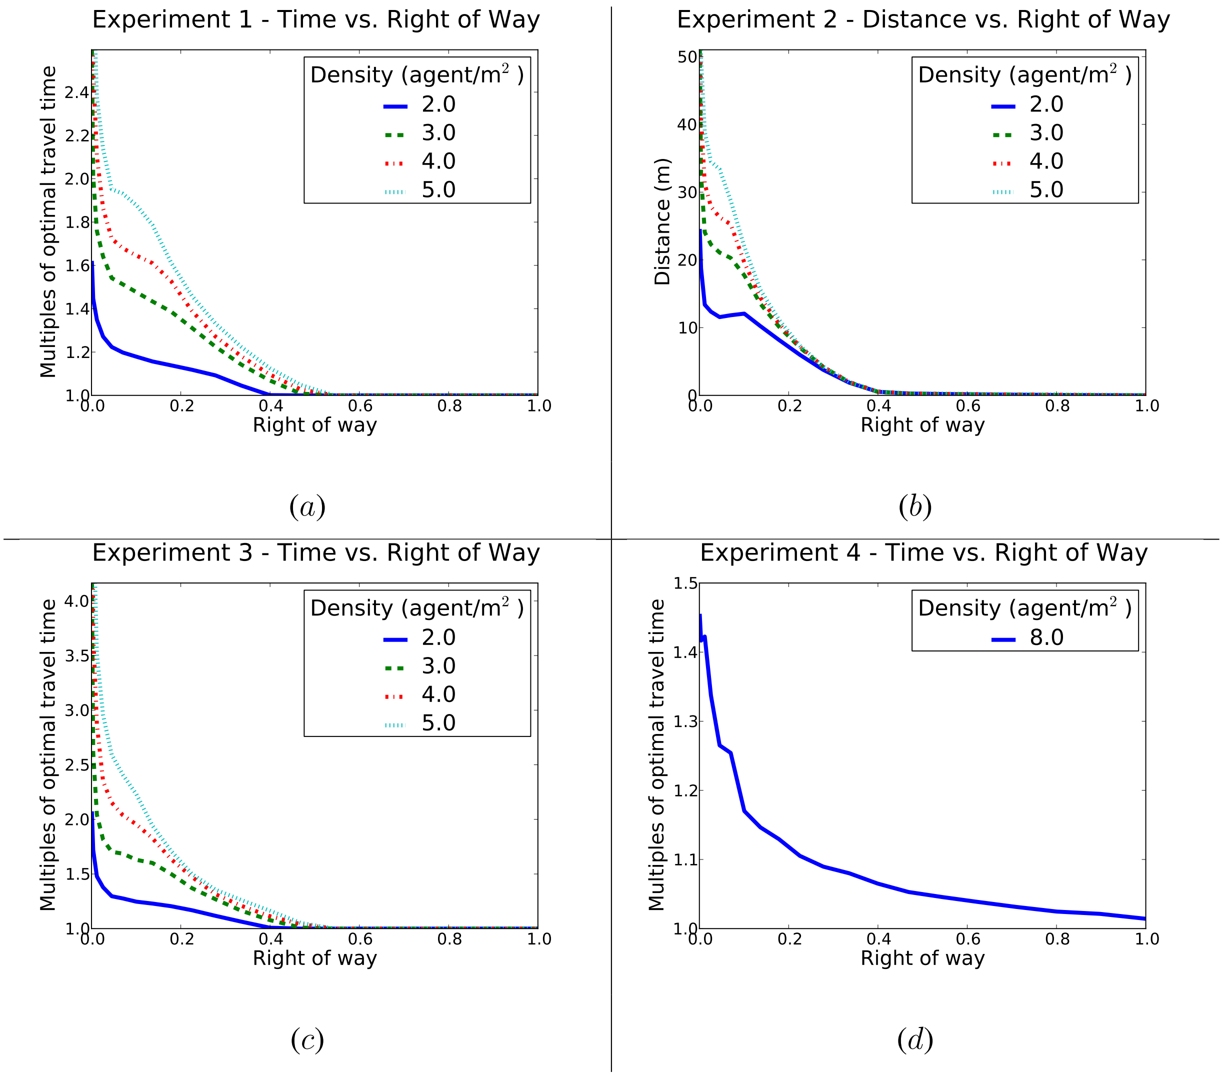

Figure 3: The impact of priority on the experiment scenarios with social forces. (a), (c), and (d) report a multiple of the baseline travel time based on right-of-way value and density. (b) shows the absolute distance traveled. Update plots with real SF analysis.

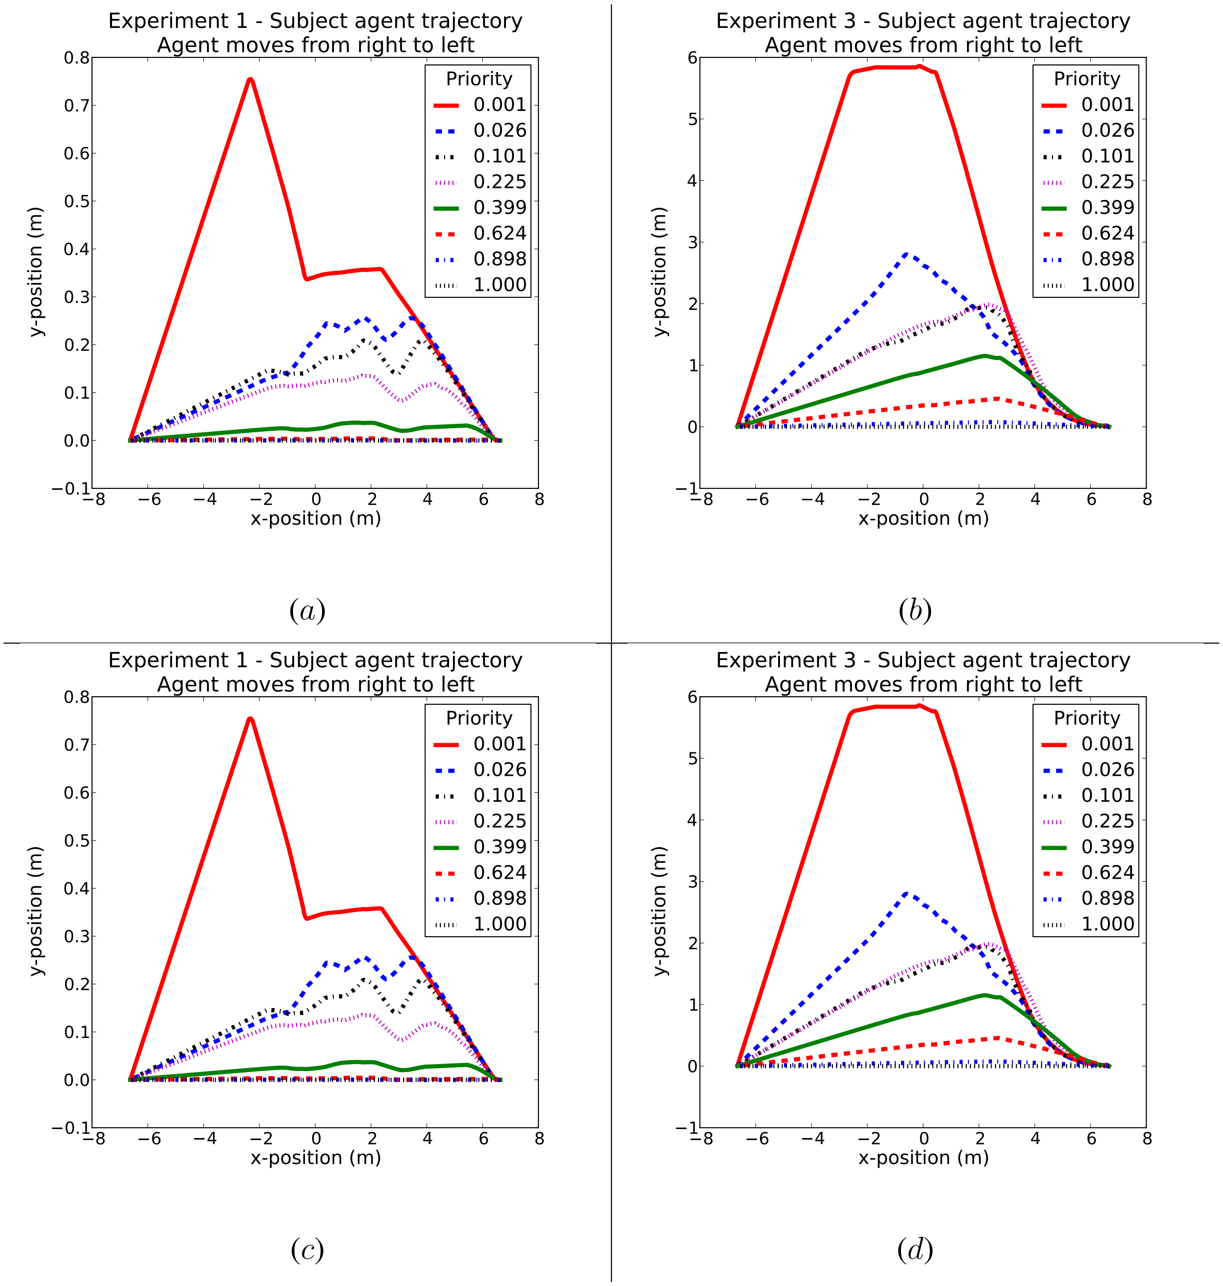

Figure 4: The trajectory of the subject agent at varying priority levels. (a) Experiment 1 (ORCA). (b) Experiment 3 (ORCA). (c) Experiment 1 (social forces). (d) Experiment 3 (social forces).

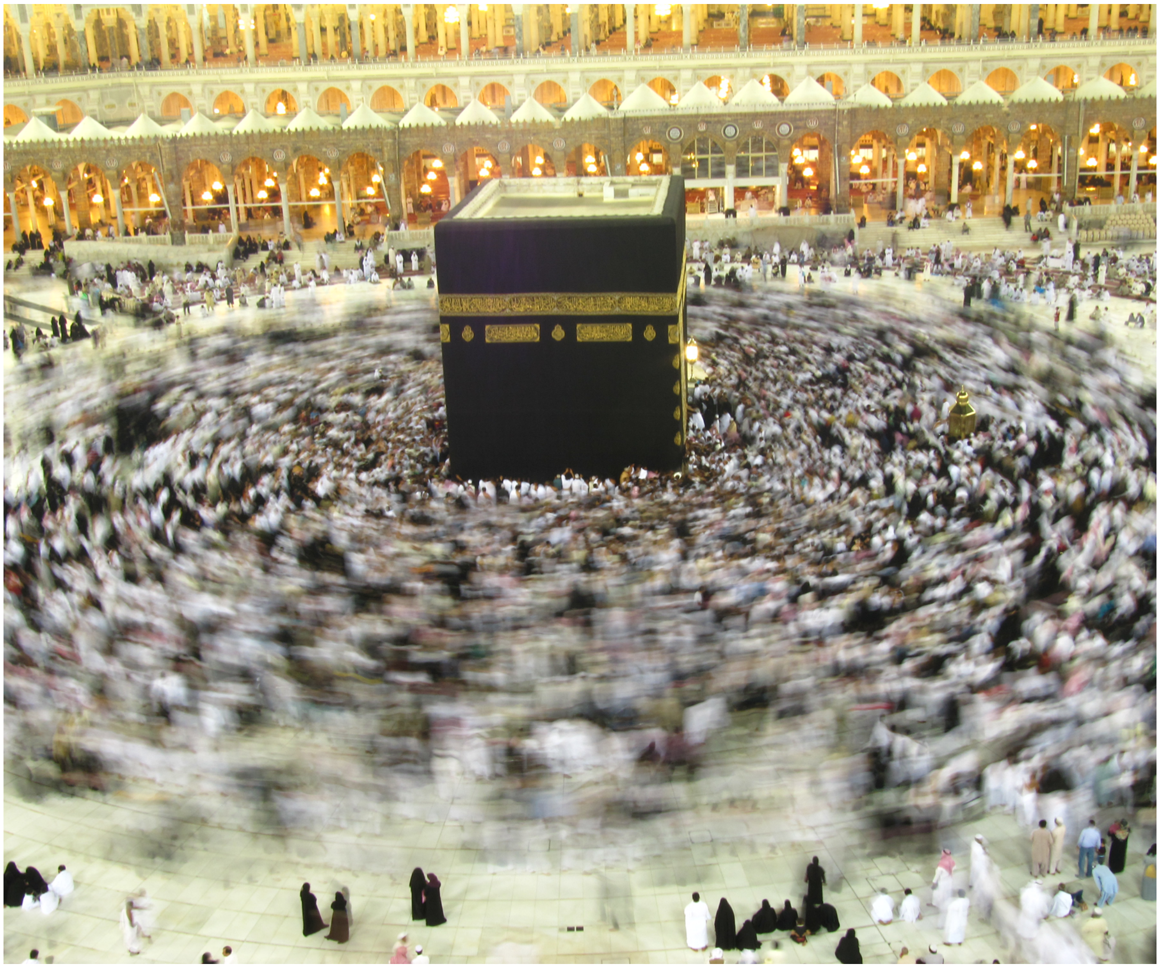

Figure 5: A photograph of pilgrims performing the Tawaf. The motion blur clearly shows the motion of the crowd. The queuing pilgrims near the wall of the Kaabah are able to maintain the queue integrity despite this flow. I'm working on obtaining permission to use this.

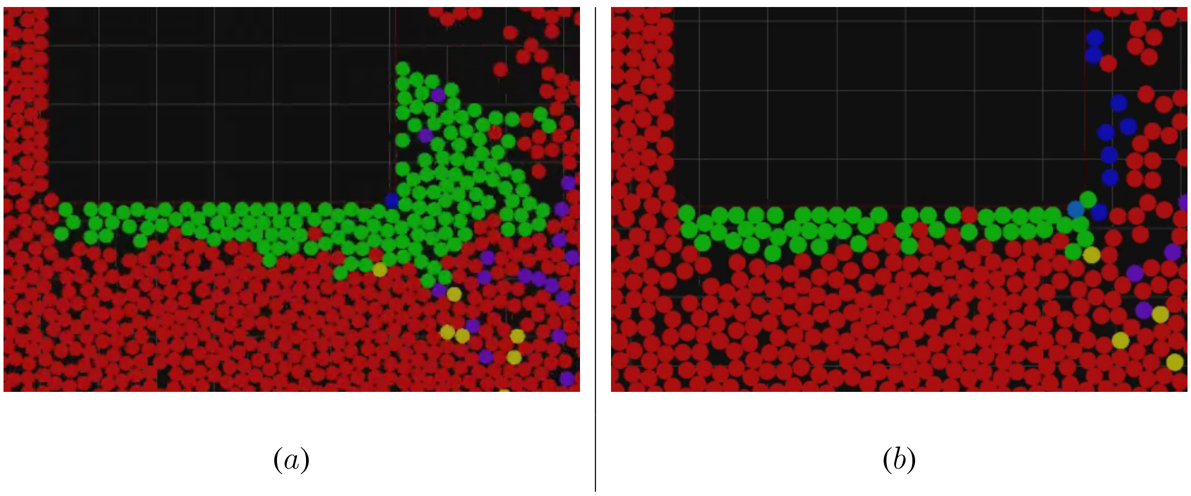

(a) (b)

Figure 6: Two frames from simulation of the Tawaf. The green agents are trying to queue. (a) These agents have no priority. The motion of the circling agents are enough to cause the queue to drift around the Kaabah, destroying the integrity of the queue. (b) These agents have higher priority and right of way. Right of way enables them to maintain the queue integrity. These will be replaced with sequences of images showing the evolution of the queue in both cases. Also, agents will be two colors, queuing agents and NON-queuing agents – and they won't be red and green.