Differential Game Theory and Pedestrian Traffic

Andrew Glaws

May 1, 2015

1 Introduction

The modeling of pedestrian traffic has been studied using a variety of techniques including fluid flow dynamics and continuum mechanics [1,2,3,4]. Here, we view pedestrian traffic as a differential game. The application of differential game theory attempts to treat pedestrians as logically-thinking individuals trying to reach their destination in the best possible manner as opposed to particles within a system whose behavior is fully determine by the overall flow.

The study of pedestrian traffic continues to be of great interest as high traffic areas, such as airports, malls, and businesses, attract larger crowds which must maneuver their environment as quickly as possible. Additionally, similar techniques to those developed for this model can be applied to other forms of traffic such as automobile traffic. The ultimate goal in modeling these traffic patterns is to determine how environmental changes can be implemented to influence more efficient traffic flow patterns from pedestrians.

This paper begins with a background discussion of optimal control theory and differential games. Then, we develop the main solution techniques which are used in this model and offer a simple one-dimensional example which illustrates how the techniques are applied. Next, the full pedestrian traffic model is developed, explained, and solved using the techniques discussed. Finally, the model is implemented in several environments and the behavior of the model is analyzed.

2 Background

2.1 Optimal Control Theory

A standard problem in optimal control theory begins with a state variable $\boldsymbol{x} \in \mathbb{R}^m$ and a control variable $\boldsymbol{u} \in \mathcal{U} \subset \mathbb{R}^n$. The set $\mathcal{U}$ is referred to as the set of admissible controls for the problem and is generally used to place practical limits on the effectiveness of the control (i.e. a rocket cannot apply an infinite amount of thrust to a shuttle). The goal of this problem is to find the optimal control, $\boldsymbol{u}^* \in \mathcal{U}$, which steers the state from some initial set $\boldsymbol{X}_0$ to a target set $\boldsymbol{X}_1$ subject to the state dynamics

$$\begin{cases} \dot{\boldsymbol{x}}(t) = \boldsymbol{f}(t, \boldsymbol{x}, \boldsymbol{u}) \\ \boldsymbol{x}(t_0) = \boldsymbol{x}_0 \in \boldsymbol{X}_0. \end{cases} \tag{1}$$

The determination of the optimal control is subject to a given cost functional

$$J(\boldsymbol{u}) = \phi\left(\boldsymbol{x}(t_1)\right) + \int_{t_0}^{t_1} L(s, \boldsymbol{x}, \boldsymbol{u})\, ds \tag{2}$$

which we attempt to minimize. This cost functional is made up of two components: the terminal cost, $\phi(\boldsymbol{x}(t_1))$, and the running cost, $L(s, \boldsymbol{x}, \boldsymbol{u})$.

The standard problem in optimal control leads to questions regarding the existence and uniqueness of an optimal control and methods for mathematically describing such a control. These questions have been studied thoroughly [5,6,7]. Some useful methods for finding optimal controls are discussed later in the section.

2.2 Differential Game Theory

A problem in differential game theory is similar to a standard problem in optimal control. However, differential games involve several players who each operate their own control. Additionally, each player is attempting to minimize their own cost functional.

An example of a two-player differential game includes system dynamics given by

$$\begin{cases} \dot{\boldsymbol{x}}(t) = \boldsymbol{f}(t, \boldsymbol{x}, \boldsymbol{u}_1, \boldsymbol{u}_2) \\ \boldsymbol{x}(t_0) = \boldsymbol{x}_0 \in \boldsymbol{X}_0 \end{cases} \tag{3}$$

and cost functionals

$$\begin{aligned} J_1(\boldsymbol{u}_1) &= \phi_1\left(\boldsymbol{x}(t_1)\right) + \int_{t_0}^{t_1} L_1(s, \boldsymbol{x}, \boldsymbol{u}_1, \boldsymbol{u}_2)\, ds \\ J_2(\boldsymbol{u}_2) &= \phi_2\left(\boldsymbol{x}(t_1)\right) + \int_{t_0}^{t_1} L_2(s, \boldsymbol{x}, \boldsymbol{u}_1, \boldsymbol{u}_2)\, ds. \end{aligned} \tag{4}$$

Hence, it becomes clear that both players can influence the dynamics of the system as well as both players' costs.

Differential games and their theory were first introduced by Rufus Isaacs [8] who primarily studied their applications in various warfare and combat scenarios. Many of the problems introduced by Isaacs were of the pursuit-evasion form, where one player attempts to capture the other. In addition to military applications, differential games have appeared in the study of space flight, political science, economics, and traffic management.

Depending on the scenario involved, differential games can be categorized based on various qualities:

- cooperative games - communication and non-conflicting goals allow the players to work together to obtain the best possible outcome
- open-loop games - the controls for each player are decided at the beginning and cannot be changed once the game begins
- sequential-move games - the players execute their controls in sequence so that later players can see what earlier players are doing
- zero-sum games - the act of improving one player's cost directly results in a worsening of another player's cost

2.3 The Maximum Principle

The calculus of variations is a well-developed tool for studying problems that arise in optimal control and differential game theory. Here, we examine one of the main results for finding optimal controls subject to a cost functional.

First, we must define a useful tool which appears in various ways across optimization, the Hamiltonian. In the case of optimal control or differential games, the Hamiltonian is given by

$$H(t, \boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{u}) = L(t, \boldsymbol{x}, \boldsymbol{u}) + \boldsymbol{\lambda}^T \boldsymbol{f}(t, \boldsymbol{x}, \boldsymbol{u}). \tag{5}$$

The vector $\boldsymbol{\lambda} \in \mathbb{R}^m$ is known as the adjoint system and is time-dependent so that $\boldsymbol{\lambda} = \boldsymbol{\lambda}(t)$. The adjoint behaves much like the Lagrangian multipliers used in optimization. Given the Hamiltonian, the state dynamics can be written in the form

$$\dot{\boldsymbol{x}} = \frac{\partial H}{\partial \boldsymbol{\lambda}} = \boldsymbol{f}(t, \boldsymbol{x}, \boldsymbol{u}). \tag{6}$$

Using results from calculus of variations, the adjoint state dynamics can be shown to be

$$\dot{\boldsymbol{\lambda}} = -\frac{\partial H}{\partial \boldsymbol{x}}. \tag{7}$$

Equations (**??**) and (**??**) result in a system of ODEs which must then be solved. Initial conditions arise from $\boldsymbol{x}(t_0) = \boldsymbol{x}_0 \in \boldsymbol{X}_0$ and some terminal conditions may be obtained from $\boldsymbol{x}(t_1) \in \boldsymbol{X}_1$. Additional conditions come from the so-called transversality condition on the adjoint system. Transversality states that

$$\begin{aligned} \boldsymbol{\lambda}(t_0) \perp \boldsymbol{T}_0 \\ \boldsymbol{\lambda}(t_1) \perp \boldsymbol{T}_1 \end{aligned} \tag{8}$$

where $\boldsymbol{T}_0$ and $\boldsymbol{T}_1$ are the tangent spaces to $\boldsymbol{X}_0$ and $\boldsymbol{X}_1$, respectively.

The main result in this section which will be used to find the optimal control is called the Pontryagin Maximum Principle, or simply the maximum principle. The maximum principle states that if $\boldsymbol{u}^* \in \mathcal{U}$ is the optimal control for a given system with associated optimal state $\boldsymbol{x}^*$ and optimal adjoint $\boldsymbol{\lambda}^*$, then $\boldsymbol{u}^*$ must minimize the Hamiltonian subject to these optimal state and adjoint variables. That is,

$$H(t, \boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{u}^*) \leq H(t, \boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{u}) \qquad \text{for any } \boldsymbol{u} \in \mathcal{U}. \tag{9}$$

In special cases, when the Hamiltonian is independent of time, it can be shown that the minimum value of the Hamiltonian is given by

$$H(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{u}^*) = 0. \tag{10}$$

This result will be important for the example that is studied in the next section.

The maximum principle allows us to determine possible optimal controls $\boldsymbol{u}^*$ by using a simple result from calculus:

$$\left.\frac{\partial H}{\partial \boldsymbol{u}}\right|_{\boldsymbol{u}=\boldsymbol{u}^*} = 0. \tag{11}$$

The optimal control is then found by combining this result from the maximum principle with the system of ODEs found above.

3 One-dimensional Example

In this section, we consider a simple one dimensional example in order to illustrate how the Hamiltonian and maximum principle are used to solve problems of optimal control. Consider the case of a single pedestrian walking down a straight hallway alone. Since only a single player is present, this is a problem of optimal control and not a differential game. The state vector is set to be the pedestrian's position and velocity

$$\boldsymbol{x}(t) = \begin{bmatrix} r(t) \\ v(t) \end{bmatrix} \tag{12}$$

and the dynamics of the system are given by

$$\dot{\boldsymbol{x}}(t) = \begin{bmatrix} \dot{r}(t) \\ \dot{v}(t) \end{bmatrix} = \begin{bmatrix} v(t) \\ u(t) \end{bmatrix}. \tag{13}$$

The only costs to minimize in this problem will be those relating to time and acceleration. We will assume a quadratic cost functional of the form

$$J(u(\cdot)) = \frac{1}{2}\int_0^{t_1} \left[u(s)^2 + 1\right]\, ds. \tag{14}$$

Notice that the final time for the problem is not given. In this case, we must solve for the final time using the state dynamics. This is called a free-time problem.

The initial and target states are given by

$$X_0 = \boldsymbol{x}(0) = \begin{bmatrix} r(0) \\ v(0) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \qquad X_1 = \left\{ \begin{bmatrix} 1 \\ \alpha \end{bmatrix} \in \mathbb{R}^2 : \alpha \in \mathbb{R} \right\}, \tag{15}$$

respectively. Hence, the pedestrian starts at the origin with zero velocity and ends at $r(t_1) = 1$ with any velocity. The Hamiltonian for the system is

$$H(t, \boldsymbol{x}(t), \boldsymbol{\lambda}(t), u(t)) = \frac{1}{2}\left(u(t)^2 + 1\right) + \lambda_1(t)v(t) + \lambda_2(t)u(t) \tag{16}$$

where $\boldsymbol{\lambda}(t) = [\lambda_1(t), \lambda_2(t)]^T$ is the adjoint system. The dynamics of the adjoint are given by

$$\begin{bmatrix} \dot\lambda_1(t) \\ \dot\lambda_2(t) \end{bmatrix} = -\frac{\partial H}{\partial \boldsymbol{x}} = \begin{bmatrix} 0 \\ -\lambda_1(t) \end{bmatrix}. \tag{17}$$

Finally, the transversality conditions of the adjoint, $\boldsymbol{\lambda}(t_0) \perp T_0$ and $\boldsymbol{\lambda}(t_1) \perp T_1$ where T_0, T_1 are tangent planes to the initial and terminal states, give us that $\lambda_2(t_1) = 0$. Thus, we obtain the system

$$\left\{ \begin{array}{ll} \dot r(t) = v(t) & r(0) = 0 \\ \dot v(t) = u(t) & v(0) = 0 \\ \dot\lambda_1(t) = 0 & r(t_1) = 1 \\ \dot\lambda_2(t) = -\lambda_1(t) & \lambda_2(t_1) = 0 \end{array} \right. \tag{18}$$

This system cannot be solved yet as $u(t)$ is still unknown. The maximum principle is used to determine that the optimal control u^* is given by

$$\left.\frac{\partial H}{\partial u}\right|_{u=u^*} = \lambda_2(t) + u^*(t) = 0 \Rightarrow u^*(t) = -\lambda_2(t). \tag{19}$$

This gives the closed differential system

$$\left\{ \begin{array}{ll} \dot r(t) = v(t) & r(0) = 0 \\ \dot v(t) = -\lambda_2(t) & v(0) = 0 \\ \dot\lambda_1(t) = 0 & r(t_1) = 1 \\ \dot\lambda_2(t) = -\lambda_1(t) & \lambda_2(t_1) = 0 \end{array} \right. \tag{20}$$

This system can be solve analytically to obtain

$$\begin{aligned} r^*(t) &= \frac{1}{2t_1^3}\left(t_1 - t\right)^3 + \frac{3}{2t_1}t - \frac{1}{2} \\ v^*(t) &= -\frac{3}{2t_1^3}\left(t_1 - t\right)^2 + \frac{3}{2t_1} \\ \lambda_1^*(t) &= -\frac{3}{t_1^3} \\ \lambda_2^*(t) &= -\frac{3}{t_1^3}\left(t_1 - t\right) \end{aligned} \tag{21}$$

To determine the optimal final time, we recognize that the Hamiltonian is time-independent so that is optimal value is 0. Hence, we have

$$H(\boldsymbol{x}^*(t_1), \boldsymbol{\lambda}^*(t_1), u^*(t_1)) = \frac{1}{2} + \left(\frac{-3}{t_1^3}\right)\left(\frac{3}{2t_1}\right) = 0. \tag{22}$$

Solving this algebraic expression for the final time yields $t_1 = \sqrt{3}$. This is then plugged into the solution to obtain

$$\begin{aligned} r^*(t) &= \frac{1}{6\sqrt{3}}\left(\sqrt{3}-t\right)^3 + \frac{\sqrt{3}}{2}t - \frac{1}{2} \\ v^*(t) &= -\frac{1}{2\sqrt{3}}\left(\sqrt{3}-t\right)^2 + \frac{\sqrt{3}}{2} \\ \lambda_1^*(t) &= -\frac{1}{\sqrt{3}} \\ \lambda_2^*(t) &= -\frac{1}{\sqrt{3}}\left(\sqrt{3}-t\right) \end{aligned} \tag{23}$$

Finally, using the result from the maximum principle, we get the optimal control to be

$$u^*(t) = -\lambda_2^*(t) = \frac{1}{\sqrt{3}}\left(\sqrt{3}-t\right). \tag{24}$$

Figure 1 shows the evolution of the state and the control in time. This result appears logical as the pedestrian attempts to balance a desire to get to the target state as quickly as possible with a desire to not accelerate too aggressively. Hence, we see a large acceleration early in order to get the velocity up. As time passes, it is less beneficial to continue to accelerate aggressively so the control diminishes to 0.

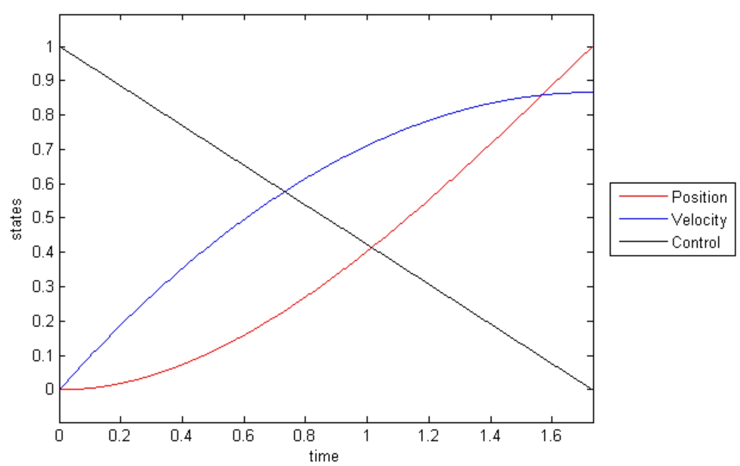

Figure 1: The evolution of the state and control variables in time for the simple one-dimension example.

4 Pedestrian Traffic Model

4.1 Building the Model

In this section, we will build up the pedestrian traffic model in the form of a differential game as described in [4]. The pedestrians are assumed to all have their own initial and target sets. As the game unfolds, the pedestrians will move based on what appears to be the best for them given their final goals and cost functional.

Several major assumptions in the model will be that

1. pedestrians continuously observe the current state and update their plan
2. pedestrians have perfect knowledge of the current state
3. pedestrians have limited prediction capabilities
4. pedestrians work harder to avoid groups that individuals

The first assumption means that we have a closed-loop control which allows the pedestrians to examine the state and decide what course of action to take at each time step. The second assumption simplifies implementation, but can be unrealistic since we know that pedestrians cannot see around corners or behind them. The third assumption results in a "temporal discount factor" which will be described later. This causes the pedestrian to more heavily weigh their cost functional in regards to events that will happen in the near future. The final assumption means that the desire to avoid other pedestrians is additive so that a pedestrian will take evasive action sooner and more dramatically to avoid a cluster of pedestrians compared to an individual.

The state of the system is made up of the locations and velocities of all the pedestrians so that

$$\boldsymbol{x}(t) = \begin{bmatrix} \boldsymbol{r}_1(t) \\ \vdots \\ \boldsymbol{r}_N(t) \\ \boldsymbol{v}_1(t) \\ \vdots \\ \boldsymbol{v}_N(t) \end{bmatrix} \tag{25}$$

where there are N pedestrians in the system. Similar to the one-dimensional example, the dynamics of the system will be given by

$$\dot{\boldsymbol{x}}(t) = \begin{bmatrix} \boldsymbol{v}_1(t) \\ \vdots \\ \boldsymbol{v}_N(t) \\ \boldsymbol{a}_1(t) \\ \vdots \\ \boldsymbol{a}_N(t) \end{bmatrix}. \tag{26}$$

The acceleration of each pedestrian, $\boldsymbol{a}_p(t)$ for $p = 1, \ldots, N$, is made up a controllable portion $\boldsymbol{u}_p(t)$ and an uncontrollable portion $\boldsymbol{w}_p(t)$. The controllable portion is obviously determined by the state and cost functional while the uncontrollable portion is used to describe physical contact with other pedestrians and boundaries. The uncontrollable portion has the form

$$\boldsymbol{w}_p = \boldsymbol{w}_{\text{bound}} + \sum_{q \neq p} \left[k_{p,1} \left((R_p^* + R_q^*) - R_{p,q} \right)^+ \boldsymbol{n}_{p,q} + k_{p,2} \left((\boldsymbol{v}_q - \boldsymbol{v}_p)^T \boldsymbol{n}_{p,q}^\perp \right) \left((R_p^* + R_q^*) - R_{p,q} \right)^+ \boldsymbol{n}_{p,q}^\perp \right]. \tag{27}$$

The $\boldsymbol{w}_{\text{bound}}$ term represents accelerations due to contact with the wall and takes a similar form to the summation, which represents accelerations due to contact with other pedestrians. The parameters $k_{p,1}$ and $k_{p,2}$ are weights for the direct contact and friction terms, respectively. The values R_p^* and R_q^* represent the physical radii of pedestrians p and q while the term $R_{p,q}$ is the distance between pedestrians p and q. Thus, if the distance between the two pedestrians is less than the sum of their physical radii, then a force is applied along the normal vector pointing from pedestrian q to p, denoted $\boldsymbol{n}_{p,q}$. In the friction term, we see that $(\boldsymbol{v}_q - \boldsymbol{v}_p)^T \boldsymbol{n}_{p,q}^\perp$ is simply the difference in the two pedestrians' velocities projected along the vector perpendicular to the normal vector described earlier. Hence, if the two pedestrians walk past each other and make physical contact, then their velocity is slowed due to a "friction" between them. As mentioned earlier, the accelerations due to contact with the boundaries behave very similarly to those due to inter-pedestrian contact.

Next, we examine the cost functional which will ultimately determine the behavior of all pedestrians in the model. It is assumed that the cost is made up of three components:

1. $L_{p,1}$: the cost of drifting from the optimal velocity
2. $L_{p,2}$: the cost of discomfort from walking too near to other pedestrians
3. $L_{p,3}$: the cost of accelerating, decelerating, or turning

Thus, the cost functional can be written in the form

$$J_p(\boldsymbol{u}_p) = \int_{t_0}^{\infty} e^{-\eta_p s} \left[c_{p,1}L_{p,1} + c_{p,2}L_{p,2} + c_{p,3}L_{p,3}\right] ds. \tag{28}$$

The exponential term enforces the third assumption mentioned earlier. The parameter η_p is the referred to as the temporal discount factor and determines how heavily to weigh issues in the near future versus those further down the road.

The cost of drifting from the optimal velocity is represented using the quadratic function

$$L_{p,1} = \frac{1}{2}\left(\boldsymbol{v}_p^* - \boldsymbol{v}_p\right)^T\left(\boldsymbol{v}_p^* - \boldsymbol{v}_p\right) \tag{29}$$

where $\boldsymbol{v}_p^*$ is the optimal velocity of pedestrian p. This velocity is determined by the shortest route (ignoring other pedestrian who may be in the way) from the pedestrian's current position to his target set. This term coaxes the pedestrian towards his final destination as quickly as possible.

The cost of discomfort from walking too near to other pedestrians can be written as

$$L_{p,2} = \sum_{q \neq p} e^{-D_{p,q}/R_p^0}. \tag{30}$$

The parameter R_p^0 is called the spatial discount factor and weighs how intensely a pedestrian works to avoid other pedestrian in relation to how far away they are from him. The value $D_{p,q}$ is a sort of distance measure between pedestrians p and q. However, the distance is weighted by whether pedestrian q is directly in front of pedestrian p or to the side. In this way, the pedestrian will work harder to avoid others who are directly in his path as opposed to those who are near him but not in his way.

The cost of accelerating, decelerating, or turning is given by the equation

$$L_{p,3} = \theta_p\left(\boldsymbol{u}_p^T\boldsymbol{e}_p\right)^2 + (1-\theta_p)\left(\boldsymbol{u}_p^T\boldsymbol{e}_p^{\perp}\right)^2 \tag{31}$$

where θ_p is a weighting parameter between accelerations longitudinally versus accelerations laterally. The normal vector $\boldsymbol{e}_p$ represents the direction that pedestrian p is facing. In this model, the pedestrian is determined to be facing the direction he is walking, that is

$$\boldsymbol{e}_p = \frac{\boldsymbol{v}_p}{||\boldsymbol{v}_p||_2}. \tag{32}$$

4.2 Solving the Model

Now that the model has been built up, we apply the techniques from earlier to solve the differential game presented. The first step is to construct the Hamiltonian

$$H(t, \boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{u}) = e^{-\eta_p t} \left[c_{p,1} L_{p,1} + c_{p,2} L_{p,2} + c_{p,3} L_{p,3} \right] + \boldsymbol{\lambda}^T \boldsymbol{f}. \tag{33}$$

In applying the maximum principle to this Hamiltonian, we first recognize that the control only appears in the third running cost functions, $L_{p,3}$, and in the state dynamics, $\boldsymbol{f}$. Hence, we have

$$\frac{\partial H}{\partial \boldsymbol{u}_p} = c_{p,3} e^{-\eta_p t} \frac{\partial L_{p,3}}{\partial \boldsymbol{u}_p} + \frac{\partial \left(\boldsymbol{\lambda}^T \boldsymbol{f} \right)}{\partial \boldsymbol{u}_p}. \tag{34}$$

Taking these derivatives and solving for the optimal control yields

$$\boldsymbol{u}_p^* = -\frac{1}{c_{p,3}} e^{-\eta_p t} M_p \boldsymbol{\lambda}_{\boldsymbol{v}_p}^* \tag{35}$$

where the matrix M_p is given by

$$M_p = \frac{1}{2} \left[\theta_p \boldsymbol{e}_p \boldsymbol{e}_p^T + (1 - \theta_p) \, \boldsymbol{e}_p^\perp \left(\boldsymbol{e}_p^\perp \right)^T \right]^{-1} \tag{36}$$

and $\boldsymbol{\lambda}_{\boldsymbol{v}_p}^*$ is the optimal adjoint associated with the optimal state velocities of pedestrian p. In order to find the optimal control, the adjoint dynamics must be solved for $\boldsymbol{\lambda}_{\boldsymbol{v}_p}^*$. We know the dynamics have the form

$$\boldsymbol{\lambda}_{\boldsymbol{v}_p} = -\frac{\partial H}{\partial \boldsymbol{x}_{\boldsymbol{v}_p}}, \tag{37}$$

and since the state dynamics, $\boldsymbol{f}$, and the running cost components, $L_{p,1}$, $L_{p,2}$, and $L_{p,3}$, are time-independent, we can write $\boldsymbol{\lambda}_{\boldsymbol{v}_p}(t, \boldsymbol{x}) = e^{-\eta_p t} \tilde{\boldsymbol{\lambda}}_{\boldsymbol{v}_p}(\boldsymbol{x})$. Hence, the adjoint system be solved to get the optimal control

$$\boldsymbol{u}_p^* = M_p \left[I - \frac{1}{\eta_p} \left(\frac{\partial \boldsymbol{v}_p^*}{\partial \boldsymbol{r}_p} \right)^T \right] \left(\frac{\boldsymbol{v}_p^* - \boldsymbol{v}_p}{\tau_p} \right) - A_p^0 M_p \left[\frac{\partial L_{p,2}}{\partial \boldsymbol{r}_p} + \eta_p \frac{\partial L_{p,2}}{\partial \boldsymbol{v}_p} \right] \tag{38}$$

where $\tau_p = (\eta_p c_{p,3})/c_{p,1}$ and $A_p^0 = c_{p,2}/(\eta_p^2 c_{p,3})$.

The first term of the optimal control attempts to keep the pedestrian from starying from the optimal velocity, $\boldsymbol{v}_p^*$. The derivative of the optimal velocity with respect to space allows the pedestrian to anticipate upcoming changes in the optimal velocity so that adjustments can be made early. The second term relates to the pedestrian's desire to avoid getting too close to other pedestrians. Finally, the costs associated with accelerating, decelerating, or turning are packaged into the constant τ_p and A_p^0 as well as the matrix M_p. These factors weigh how heavily accelerations in the two terms are applied.

5 Results

For the first problem, we examine a straight hallway with pedestrians walking from one end to the other. The code is such that at each time step, the pedestrians observe the current state of the system and update their control. The optimal velocity is constant in this case since the pedestrians are only concerned with reaching the other end of the hallway. Thus, the optimal control simplifies to

$$\boldsymbol{u}_p^* = M_p \left(\frac{\boldsymbol{v}_p^* - \boldsymbol{v}_p}{\tau_p} \right) - A_p^0 M_p \left[\frac{\partial L_{p,2}}{\partial \boldsymbol{r}_p} + \eta_p \frac{\partial L_{p,2}}{\partial \boldsymbol{v}_p} \right]. \tag{39}$$

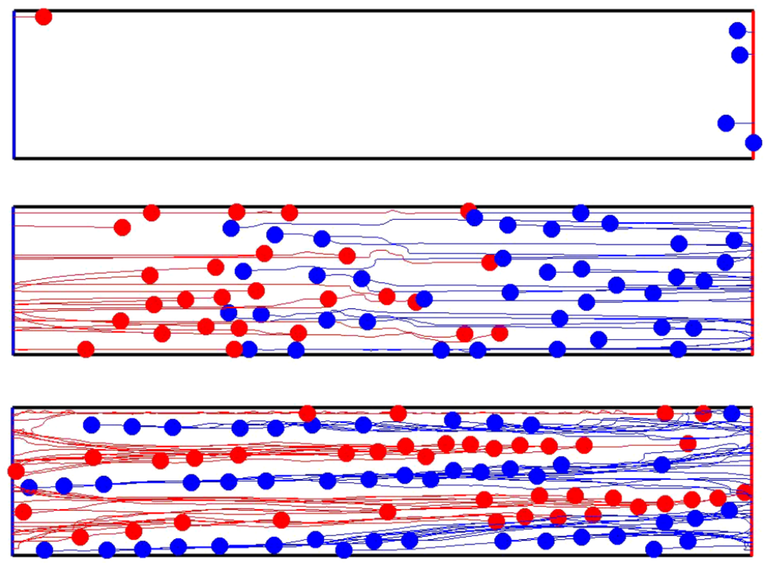

Figure 2: The evolution of pedestrian flow for a straight hallway. The target set of the pedestrians are indiciated by their color. In the long term, we see lanes begin to naturally form.

As the model unfolds, we see lanes form which allow pedestrians to avoid any direct interactions with other pedestrians and reach their final destination quickly. Figure 2 shows the overall flow as the individuals progress through the hallway. In addition to the formation of lanes, we see that pedestrians along the walls have a greater difficulty reaching their destination than those in the middle of the hallway.

For the second problem, we examine a four-way intersection with pedestrians coming from two directions and attempting to walk straight through the intersection. The optimal velocity is constant within their straight hallway for the pedestrians. However, if inter-pedestrian interaction forces a pedestrian from their straight hallway, then the optimal velocity guides them back into the flow of traffic heading to their same destination.

Similar to the formation of lanes in the straight hallway, we see the eventual formation of diagonal groups of pedestrians heading in each direction in Figure 3. This pattern allows the pedestrians to move past each other in optimal time by inter-weaving their walking patterns. Additionally, we begin to see some build-up at the corner as pedestrians who were pushed out their straight hallways attempt to return. This phenomenon is examined further in the last example.

The final problem is again a four-way intersection. However, now pedestrians are coming from all directions and heading towards any other random exit. This allows for interesting examination of traffic patterns near the corners. We see heavy build-ups at these location which are not quickly resolved as pedestrians compete for position. This appears to be one of the major flaws within the model. Despite corners being a potentially difficult part of real-life pedestrian traffic, we generally do not see the kind of build-ups shown in Figure 4. This is discussed further in the next section.

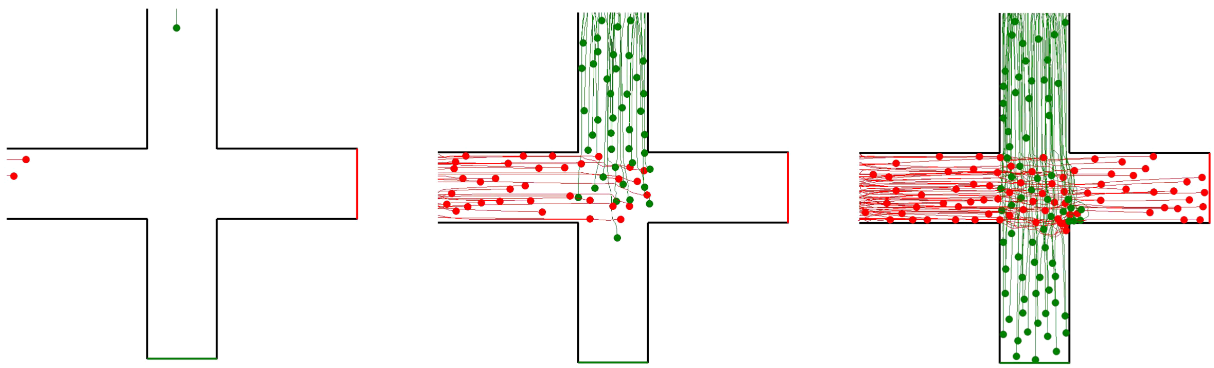

Figure 3: The evolution of pedestrian flow for a four-way intersection. The target set of the pedestrians are indiciated by their color. In the long term, we see diagonal groups begin to naturally form in the intersection area.

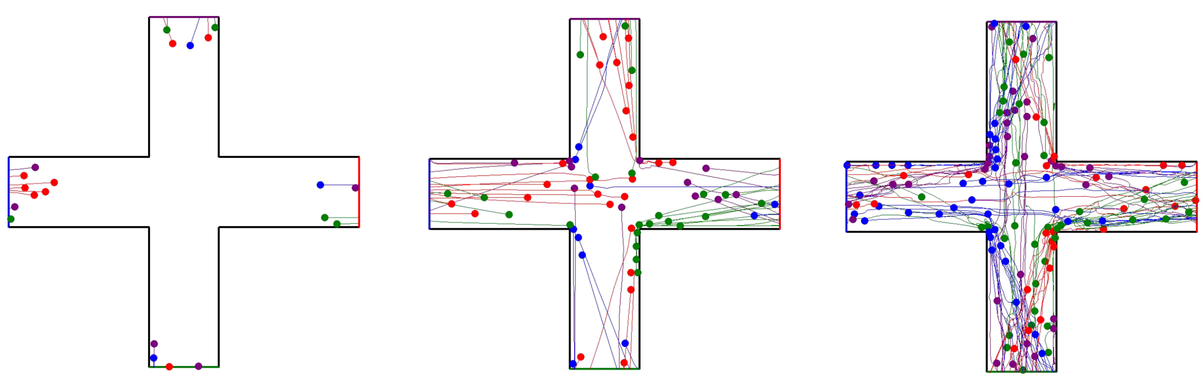

Figure 4: The evolution of pedestrian flow for a four-way intersection. The target set of the pedestrians are indiciated by their color. This problem highlights the models difficulty in handling behavior near corners where we see heavy build-ups that take a long time to dissipate

6 Conclusions

Overall, the models performance closely resembles the style and form that we expect from crowds of pedestrians. The major qualitative behavior which we observe is the formation of lanes or groups by pedestrians in order to more quickly and conveniently reach their destination. The model does have difficulties associated with pedestrians attempting to turn corners. However, this is commonly an area of difficulty for pedestrians in real life who cannot see what is coming around the corner.

Several additional implementations may be added to the model in order to more realistically represent the flow of pedestrians. The first is to add a level of cooperation to the game. In real life, pedestrians' goals are generally non-conflicting and communication (verbal or nonverbal) is not restricted. Thus, we would expect a level of cooperation to improve the flow of traffic, especially at difficult junctions such as corners. Implementing cooperation can be done by applying the maximum principle to all the controls at once such that

$$\left.\frac{\partial H}{\partial \boldsymbol{u}}\right|_{\boldsymbol{u}=\boldsymbol{u}^*} = \mathbf{0}. \tag{40}$$

In this way, pedestrian p has some understanding of how the other pedestrians will likely update their controls for the next time step and can plan accordingly. The second addition to the code would be to limit the pedestrians' knowledge of the state. In real life, pedestrians cannot see behind them, around corners, or around other pedestrians, and this fact can complicate things. This can be done by applying some observation matrix to the current state in order to obtain the observable state

$$\boldsymbol{y}_p = B_p \boldsymbol{x}. \tag{41}$$

The same solution techniques are then applied to the observable state rather than the entire state.

Lastly, the model contains a variety of seemingly arbitrary parameters such as

- R_p^0 - spatial discount factor
- η_p - temporal discount factor
- $k_{p,1}$
- $k_{p,2}$ } weighting factors for inter-pedestrian interactions
- $c_{p,1}$
- $c_{p,2}$ } weighting factors for costs functions
- $c_{p,3}$
- c_p^+
- c_p^- } weighting factors for longitudinal/latitudinal discomfort
- θ_p - relative weight of longitudinal vs. latitudinal acceleration

A deeper understanding of how these parameters affect the flow of pedestrians needs to be obtained in order to build the best possible model. This can be done by studying pedestrian flow either on a macroscopic scale (general flow patterns) or a microscopic scale (individual behaviors) and using this data to determine the best parameter values. Without understanding the true impact of these parameters on the model, the results cannot be practically applied to real-life problems.

References

1. Hoogendoorn, S. P., Bovy, P. H. L. Simulation of pedestrian flows by optimal control and differential games. *Optimal Control Applied Methods*. 24:153-172. 2003.

2. Helbing, D. A fluid-dynamic model for the movement of pedestrians. *Complex Systems*. 6: 391-415. 1992.

3. Hughes, R. L. A continuum theory for the flow of pedestrians. *Transportation Research Part B*. 36:507-535. 2002.

4. Hoogendoorn, S. P., Bovy, P. H. L. Pedestrian route-choice and activity scheduling theory and models. *Transportation Reasearch Part B*. 38:169-190. 2004.

5. Burns, J. A. *The Calculus of Variations and Control with Modern Applications*. Boca Raton: CRC. 2014.

6. Evans, L. C. An Introduction to Mathematical Optimal Control Theory. 2nd Ed. University of California, Berkeley.

7. Johnson, P. A. Numerical Solution Methods for Differential Game Problems (Master's thesis). Massachusetts Institute of Technology, Cambridge, MA. 2007.

8. Isaacs, R. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. New York: Wiley. 1965.