




Multi-agent Path Planning with Heterogenous Interactions in Tight Spaces

V. Modi,¹  Y. Chen,¹ A. Madan,¹ S. Sueda² and D. I. W. Levin¹

¹University of Toronto, Toronto, Canada
vismay@cs.toronto.edu, cheniyixin1008@gmail.com, {amadan, diwlevin}@cs.toronto.edu
²Texas A&M University, College Station, TX, USA
sueda@tamu.edu

Abstract

By starting with the assumption that motion is fundamentally a decision making problem, we use the world-line concept from Special Relativity as the inspiration for a novel multi-agent path planning method. We have identified a particular set of problems that have so far been overlooked by previous works. We present our solution for the global path planning problem for each agent and ensure smooth local collision avoidance for each pair of agents in the scene. We accomplish this by modelling the collision-free trajectories of the agents through 2D space and time as rods in 3D. We obtain smooth trajectories by solving a non-linear optimization problem with a quasi-Newton interior point solver, initializing the solver with a non-intersecting configuration from a modified Dijkstra's algorithm. This space-time formulation allows us to simulate previously ignored phenomena such as highly heterogeneous interactions in very constrained environments. It also provides a solution for scenes with unnaturally symmetric agent alignments without the need for jittering agent positions or velocities.

Keywords: motion planning, animation

CCS Concepts: • Computing methodologies → Computing methodologies → Modeling and simulation

1. Introduction

On a hot dry day in Kruger National Park, an empty truck idles on the side of a road. Sam, the driver of the truck, has wandered a hundred metres off the road in an attempt to take a picture of a tree with 10 baboons. The baboons suddenly jump out of the tree and charge towards Sam. Sam panics and starts running back to the truck; however, coincidentally an elephant wandering in the area is on a path perpendicular to Sam's, between her and the truck. How will Sam get back to the truck safely while outrunning the baboons and avoiding collision with the elephant? What are the paths of the 12 agents in the scene: Sam, the elephant and the 10 baboons?

This problem illustrated in Figure 1 above poses many challenges to a multi-agent path planning algorithm. First, Sam must anticipate collisions ahead of time, in order to move quickly and efficiently to their truck. Second, the three groups of agents, Sam, the baboons and the elephant, have dramatically different masses and behaviours. For instance, while Sam seeks to avoid all animals, the massive elephant is untroubled, and will stubbornly continue on its path. Finally, geographic features such as additional trees and

ponds can lead to a highly constrained environment. Existing state-of-the-art crowd simulation methods struggle to compute anticipatory agent paths in constrained environments when heterogeneous interactions are involved (Table 1) making them ill-suited for application in planning problems such as the example given above.

We propose a new multi-agent path planning algorithm well-suited for these problems. Our model directly optimizes the space-time trajectories of all agents which allows for per-agent physical and psychological characteristics and smooth anticipatory trajectories. A novel, differentiable space-time repulsive energy ensures collision-free trajectories. Using our approach, Sam arrives safely at the truck before the baboons.

2. Related Work

Successful multi-agent path planning requires an algorithm to both correctly model the behaviour of independent agents as well as their interactions. A common approach is to apply a dynamics model based on Newton's second law of motion which is integrated over

Table 1: Comparison of (1) ORCA [VDBGLM11], (2) Implicit Crowds [KSNG17], the brand new time-to-collision method (3) NH-TTC [DKG20], (4) Repulsive Curves [YSC21], (5) Continuum Crowds [TCP06], (6) Space-Time Planning With Parameterized Locomotion Control [LLKP11] and (7) Modular Framework for Adaptive Agent Base Steering ([SKH*11]). Y—feature is available, N—feature is not possible, P—feature might be possible, but not demonstrated.

Features	Ours	(1) ORCA	(2) IC	(3) NH-TTC	(4) RC	(5) CC	(6) STPL	(7) AABS
Asymmetric interactions	Y	N	P	P	N	N	N	N
Extremely constrained environments	Y	N	N	P	Y	N	P	Y
Smooth local interactions	Y	N	Y	N	Y	N	Y	N
Intuitive control parameters	Y	Y	N	P	N	N	N	Y
Multiple agents	Y	Y	Y	Y	Y	Y	N	Y
Code available online	Y	Y	Y	Y	Y	N	N	N




Figure 1: Sam needs to escape, from the baboons while avoiding the elephant, whereas the elephant is unconcerned with the other agents in the scene and only avoids the pond.

time to produce plausible agent trajectories. Psychological and social characteristics of the group can be incorporated into the dynamics equations as Social Forces (SF) [PGOSB05, HM98, WLJT17]. Inter-agent forces, such as those that handle collision avoidance are added through a variety of means, and we can partition the space of successful approaches based on locality of their models in both space and time. Approaches such as Guy *et al.* [GKLM11] incorporate psychological factors into an underlying dynamics model, but it is impossible to tell apart the psychological traits of the agents by simply observing the simulation. Our method makes the impact of agent characteristics on the trajectory very obvious, while providing a standalone local and global dynamics model for the scene.

In local methods, agent decision making requires information only local in space and in time. Local collision avoidance methods [VDBGLM11, GNCL14, WLP16] compute collision response using local information in space and time (the planning horizon can be as small as a single time step). These methods struggle to generate smooth anticipatory collision responses. Optimization-based methods such as Implicit Crowds [KSNG17] and Crowd Patches [RLC*14, YMPT09] attempts to overcome this difficulty. Implicit Crowds introduces a time-to-collision potential to the crowd dynamics. This gives agents richer space-time information on which to act; however, this energy is still effectively local in time, computed from the current state of the system (position and velocity). Crowd Patches optimizes for collision-free trajectories

in highly local, patches with fixed start and end positions in space and time. NH-TTC [DKG20] improves upon prior work by utilizing longer planning horizon, geometrically represented by curves in space. Intersection checks between these curves allow agents to react to collisions likely to occur in the near future. Similarly, vision-based methods, such as Refs. [OPOD10, DMCN*17] provide an anticipatory collision avoidance model, but with limited path planning and anticipatory motions and potentially jagged, highly non-smooth, bumpy agent motion. Data-driven approaches such as Charalambous and Chrysanthou [CC14] use an underlying state-action graph created via external data to generate trajectories. However, these trajectories are highly data-dependent, do not factor in environmental constraints, only seem to operate in sparse crowds and must be used in conjunction with some other higher level global path planner.

An alternative approach to more local methods is to extend the collision response globally in space using a fluid like pressure solve [Hug02, TCP06, NGCL09]. However, because these methods only consider the configuration of the system (position and velocity of each agent) at a single time, their ability to produce smooth anticipatory collision responses, especially in a sparser setting, is reduced. Continuum Crowds [TCP06] stands out as one of the only methods that incorporates both a global planner through Dijkstra’s search and local collision avoidance through a fluid-like pressure solve. Another option is to handle local collision avoidance using Optimal Reciprocal Collision Avoidance (ORCA) or SF, as done by the data-driven method [TYK*09], as well as the video-based method [Sta14] (uses the RVO2 library’s implementation of ORCA) and the Integer-Linear-Programming-based method [KGvdS13] (uses SF). One might use [HKHL13], another data-driven method which uses a *purely* stochastic collision avoidance method. Another video-based crowd synthesizer, [FR13], develops the idea of using *crowd-tubes* (a set of video-generated space-time trajectories) which are used to synthesize novel scenes from video-clip data. However, this limits scene creation to looping repeatedly tiled clips of existing video data in wide-open environments. Importantly, none of the aforementioned approaches solve the problems of handling tight environmental constraints or heterogeneous agent interactions for many reasons such as a lack of global planning in some methods, an inability to anticipate motion in others and homogeneous agent interactions in most. In fact, an adjacent field of research that involves measuring the ‘correctness’ of various crowd models [GVDBL*12, WJGO*14] also




Figure 2: The repulsive curves energy forces agents to be maximally far away from one another which is not an intuitive behaviour.

fails to test for heterogeneous agents and behaviour in complex environments.

Finally, while not strictly designed for multi-agent path planning, Repulsive Curves [YSC21] can potentially be used for agent planning. While it yields impressive results in 3D curve untangling, when applied to multi-agent path planning, it has several drawbacks. First, in unconstrained environments, Repulsive Curves will maximally repel agents away as shown in Figure 2. Second, like other previous methods, there is no straightforward way to model asymmetric agents. Last, Repulsive Curves use random ‘jitter’ to ensure that no trajectories initially overlap. From our own experiments, we have observed that this is not sufficient to guarantee non-overlapping initial trajectories. Meanwhile, our space–time Dijkstra’s approach is guaranteed to create initially non-intersecting trajectories, but is not sufficient for overcoming the other limitations of Repulsive Curves.

An attempt to create a cohesive short and long-term path planning algorithm, van Toll and Pettré [vTP19] introduce the concept of ‘Navigation Strategies’. The short-term collision avoidance can be replaced by any existing algorithm (ORCA, SF, *etc.*), same as the long-term path planner (Dijkstra, A*, *etc.*). The strategic layer exchanges information from both, re-running the path planner when it is apparent that the collision avoider has encountered an obstacle. This approach creates an interface to balance the short- and long-term incentives of an agent; however, agents are not guaranteed to reach the goal, and the approach is likely to fail in very tight scenes. Furthermore, it does not actually resolve the underlying issues of the short-term collision avoidance methods, but attempts to use long-term planning to avoid those issues.

In this paper, we present an alternative long- and short-term multi-agent path planner that computes its response *globally in space and time*. Previously, space–time trajectory optimization [WK88] was initially applied to keyframe interpolation, with subsequent methods such as [PSE*00] allowing manual editing of object trajectories and [SKF08] allowing interactive motion correction and synthesis using graph search methods. The concept of using space–time graph search methods for collision avoidance is furthered in Ref. [LLKP11] for a single agent and in Ref. [SKH*11] for small groups. For large scale situations, standard space–time approaches would be computationally prohibitive, but [KGvdS13] introduce the idea of using network flow optimization over space–time for performance. Path planning in the robotics world explores many different

topics from time-optimal end-effector manipulation [RGHD13] to gait trajectory optimization [WBHB18] to graph-based multi-agent path planning [YL16]. The graph-based path planner described in Ref. [vdBO04] creates pre-planned roadmaps for agents to move through space and time and [vdBO05] establish an order of priority within agents. The drawback of these path planners is that they only work in constrained spaces, unlike our method which works in both wide open spaces as well as tightly constrained spaces. Our method builds on all these prior methods by expanding the look-ahead globally and ensuring smooth paths by collision resolution on the entire trajectories rather than using a limited space–time graph search approach. As explained in the next section, our agents exhibit an implicit prioritization as well, but its effects are mitigated by the subsequent optimization step.

Our method treats each agent as an individual space–time curve with only three requirements: a starting position, a starting time and an ending position (an exact end time is not required). We use a globally supported, differentiable LogSumExp smooth distance in space–time to guarantee collision-free trajectories and solve the resulting problem using an interior-point technique. While we do not claim to perfectly account for human-like constraints, our agents do exhibit anticipatory behaviours such as slowing and waiting to let others pass. This method makes no assumptions about agents having identical mass or other physical properties. This enables planning of intricate scenes with environmental constraints such as the one described in the introduction.

For small-to-medium scale crowds, our method outperforms current state-of-the-art methods in simple scenarios (Figures 3 and 4) and more complicated, constrained environments. ORCA, NH-TTC and Implicit Crowds were chosen as three well-known contemporary methods with author-provided open source implementations. Additionally, we use the SteerSuite [SKFR09] implementations for Plan–Predict–React (PPR) and SF for our comparisons. For all methods in the comparisons, we keep timestep at 0.1 s, the agent placement is standardized and additional parameters for each method can be found in Refs. [NX15a, NX15c, NX15b, Unk15, Dav20]. The other methods fail in Figure 4 due to being unable to handle the symmetric nature of the problem (agents push against each other with equal forces and get stuck), or they produce an unnaturally choreographed trajectory in Figure 3. Our method, on the other hand, does not have a timestep, does not suffer from the failure under perfect symmetry, and produces results most comparable to the real-life human trajectories shown in the figures.

3. Method

At a high level, we are influenced by the notion of correlated equilibrium described in Ref. [Aum87]. In a correlated equilibrium solution to a non-cooperative game, an ‘oracle’ chooses a strategy for each player, and no player has any reason to deviate from the chosen strategy assuming others do not deviate either. The result is an equilibrium solution which maximizes collective utility.

Our method acts as the oracle of the scene and plans agent trajectories in a way that collectively maximizes the utility of the entire scene. Additionally, unlike previous approaches which ‘pre-set’ trajectories for agents (can be done artistically as in Ref. [CVTZ*22],

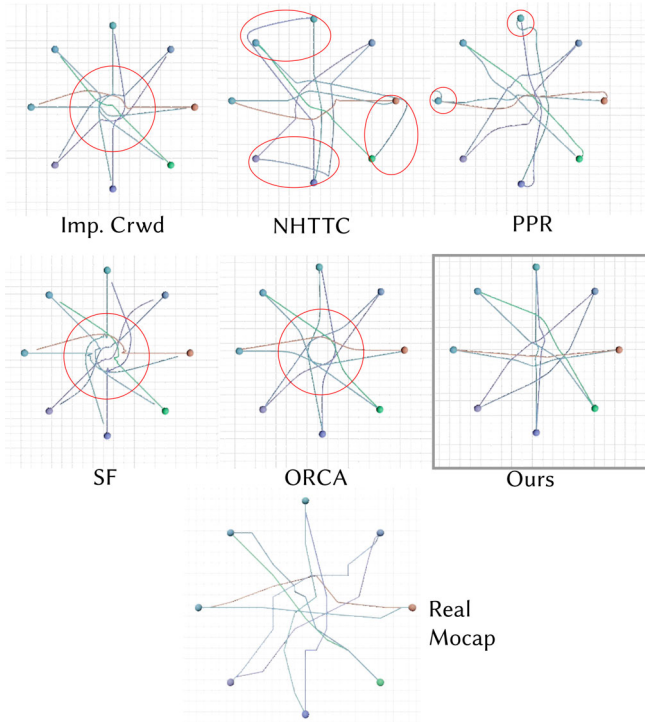


Figure 3: We show comparisons with several other methods for a scene with eight agents crossing the diameter of a circle. Our method performs most similar to the real motion-captured trajectories without any of the strange or unnaturally choreographed behaviour of other methods. The red circles highlight overly choreographed, overly symmetric and unnatural motion in prior methods.

or with navigation meshes as in [vTTK*16]), we allow our agents to find their own utility-maximizing trajectories. Last, real-life agents have different sizes, masses and personalities, which affect the agents’ utilities, and therefore, their paths as well. Our local–global path planning approach allows for these nuances.

3.1. Paths map to costs

As shown in Figure 5, for an agent, each path from the start location to the end location through space maps to some scalar cost value. For example, if the agent’s cost function minimizes distance travelled, a straight line from start to end would prove to be the optimal path. In Figure 5, the 2D x - y plane describes the domain of spatial motion for agent paths. Agents do not simply move through space, but also move through time which is denoted by the vertical z -axis. So, in 3D space–time, the agent’s motion through space (x, y) and time (t) is described by a 3D curve which corresponds to a cost. For any given moment in time, the projection of the curve onto the x - y plane gives us the agent’s location.

Let us consider a simple case where there are no other agents in the scene and the terrain is flat and free of obstacles. The agent, indexed henceforth by subscript a , wishes to minimize its cost, $\tilde{\Psi}$; the tilde symbol indicates the variable or function is continuous, not discrete. Lower case variables indicate a single agent. Upper

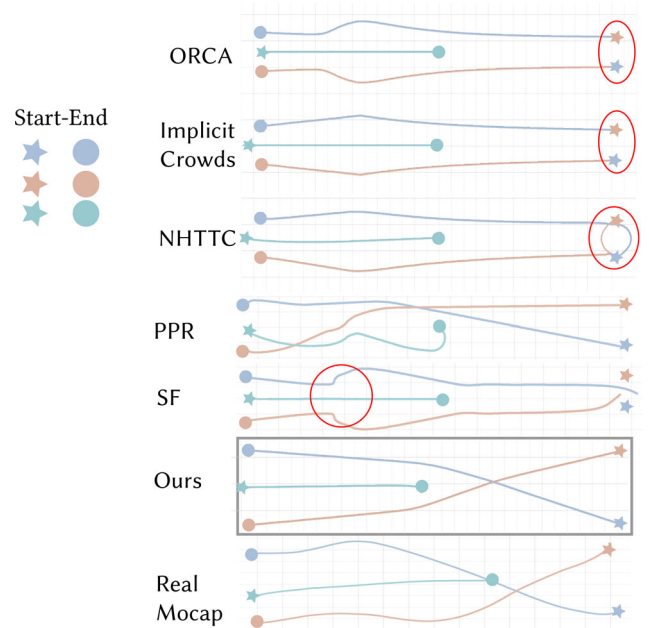


Figure 4: We show comparisons for a simple scene with three agents. Only our method lets agents reach actual their desired ends (denoted by stars) smoothly. The red circles highlight the unintuitive effects of several prior method.




Figure 5: Top: A path through space can be represented as a curve embedded in R^2 . Every path has a scalar cost (utility) value. For a cost function that minimizes distance travelled, f_1 and f_2 are sub-optimal paths from the start to the goal (red X), but f_3 is optimal. Bottom: A path through space and time can be represented as a curve in R^3 . A space–time rod (as shown here) is simply a 3D curve with a collision radius.




Figure 6: This space–time rod describes a discrete space–time trajectory curve made up of many nodes. Our agent moves along the projection of the rod from its start to end, taking into account the boundary constraints of the problem.

case indicates that the variable aggregates all agents. The motion of the agent through space and time denoted by the agent's 3D space–time curve $\tilde{\mathbf{s}}_a$ embedded in $R^{x,y,t}$ requires constraints. We linearize all our constraints and lump them together into \tilde{B} . Put together, the full optimization for a single agent a is

$$f_a^* = \min \int_{\tilde{\mathbf{s}}_a} \tilde{\psi}(\tilde{\mathbf{s}}_a) d\tilde{\mathbf{s}}_a \quad (1a)$$

$$s.t. \tilde{B}\tilde{\mathbf{s}}_a \leq \mathbf{0}. \quad (1b)$$

The cost, Equation (1a), for agent a is integrated over the trajectory of the agent ($\tilde{\mathbf{s}}_a$). The specific nature of the cost function determines the agent's behaviour based on its characteristics. An agent might have a preferred walking speed, or a stubbornness factor, or a radius of comfort, all of which (and more) can be encoded into components of $\tilde{\psi}$. In order to solve this optimization, we must first discretize the agent's trajectory into the discrete curve \mathbf{s}_a shown in Figure 6. We also discretize the cost $\tilde{\psi}$ into separate intra-agent costs, which solely affect the path of one individual, and interaction costs, which can affect multiple agents.

3.2. Discretizing

Our agent's cost function takes in the agent's **discrete** 3D space–time curve \mathbf{s}_a as input and outputs a scalar cost for that path, f_a . We discretize the agent's continuous $\tilde{\mathbf{s}}$ into a piecewise linear curve described by $n + 1$ nodes $\mathbf{s}_a = [(x_a^0, y_a^0, t_a^0), \dots, (x_a^i, y_a^i, t_a^i), \dots, (x_a^n, y_a^n, t_a^n)]$ for nodes $i = 0..n$ connected sequentially by edges. We must also discretize our constraints. First, even though time is a variable in our formulation, Equation (2b) ensures that the agent cannot move backwards in time. Second, the agent has a start location (x_a^0, y_a^0) and start time t_a^0 denoted in Equation (2c). Third, Equation (2d) sets a goal or an end location (x_a^n, y_a^n) . Last, the agent cannot take an infinite amount of time, so Equation (2e) bounds the agent by a max time T_a^{max} .

Putting all of this together, we can re-write the optimization for agent a with the discrete generalized cost function ψ as

$$f_a^* = \min \sum_{i=0}^n \psi(\mathbf{s}_a) \quad (2a)$$

$$s.t. t_a^i \leq t_a^{i+1} \quad (2b)$$

$$(x_a^0, y_a^0, t_a^0) = (\mathbf{x}_a, \mathbf{y}_a, \mathbf{t}_a)^{start} \quad (2c)$$

$$(x_a^n, y_a^n) = (\mathbf{x}_a, \mathbf{y}_a)^{end} \quad (2d)$$

$$t_a^n \leq T_a^{max}. \quad (2e)$$

Now with a template for our optimization problem, we can replace the generalized cost function with specific discrete costs. Our specific cost functions are derived from observation of real behaviour. These behavioural observations are divided into two categories, intra-agent costs and interactions costs. Intra-agent cost functions only look at one agent's path at a time. Interaction cost functions include avoiding collisions with other agents, collisions with static obstacles such as walls or furniture, gathering behaviour within friends, heterogeneous behaviour based on the mass, size of the agents or other characteristics.

3.3. Intra-agent costs

Since these costs apply to a single agent, we calculate intra-agent costs for arbitrary agent a and later we show how to sum the costs over the entire scene. The path of our agent, \mathbf{s}_a , is comprised of $n + 1$ nodes indexed by $i = 1..n$. Each node $(x^i(\mathbf{s}_a), y^i(\mathbf{s}_a), t^i(\mathbf{s}_a))$ is comprised of the agent's spatial (x, y) coordinates and time (t) coordinates. For all cost functions, agent a also has constant weighting terms, K_a , as well as constant characteristics such as mass, m_a , and preferred end time, T_a^p .

3.3.1. Intra-agent kinetic cost

Our agent a will try to minimize energy expenditure by minimizing the action over time. This model is similar to the Principle of Least Effort proposed in Ref. [GCC*10], except rather than minimizing it in a local-greedy fashion, we minimize effort over the entire trajectory. In our kinetic cost model, the rest-state expenditure is assumed to be negligible (agents are snacking, grazing or powered off), and only the kinetic expenditure is considered. Given an agent's space–time curve \mathbf{s}_a , kinetic cost of the agent over the curve can be written as

$$C_a^K(\mathbf{s}_a, K_a^K) = K_a^K \sum_{i=0}^n \frac{1}{2} m_a \frac{(\Delta \mathbf{x}_a^i)^T (\Delta \mathbf{x}_a^i)}{(\Delta t_a^i)^2} (\Delta t_a^i) \quad (3)$$

$$= K_a^K \sum_{i=0}^n \frac{1}{2} m_a \frac{(x^{i+1}(\mathbf{s}_a) - x^i(\mathbf{s}_a))^2 + (y^{i+1}(\mathbf{s}_a) - y^i(\mathbf{s}_a))^2}{t^{i+1}(\mathbf{s}_a) - t^i(\mathbf{s}_a)} \quad (4)$$

where Equation (3) is the kinetic energy $\frac{1}{2}mv^2$ for curve segment $e_a^i = [x_a^i, y_a^i, t_a^i, x_a^{i+1}, y_a^{i+1}, t_a^{i+1}]$ integrated over total time travelled $\Delta t_a^i = t_a^{i+1} - t_a^i$. This simplifies into Equation (4) where K_a^K is the agent-wise weighting coefficient, m_a is the mass of the agent and s_a is the discretized path curve (our input variable) made up of $n + 1$ nodes.

3.3.2. Intra-agent acceleration cost

In order to penalize acceleration in agent a 's trajectory, the acceleration cost

$$C_a^A(s_a, K_a^A) = K_a^A \sum_{i=1}^{n-1} \frac{1}{2} (\theta_a^i)^2 \quad (5)$$

measures the curvature of s_a through the discretized acceleration cost where angle θ is the angle between two piece-wise linear segments of the curve computed using the stable arctan function $\arctan2$. The angle is $\theta_a^i = \arctan2\left(\frac{\|(s_a^{i+1} - s_a^i) \times (s_a^i - s_a^{i-1})\|}{(s_a^{i+1} - s_a^i)^T (s_a^i - s_a^{i-1})}\right)$ where $s_a^i = [x^i, y^i, t^i]$ for each node in the curve.

3.3.3. Intra-agent preferred end time cost

Any agent a has a preferred end time, T_a^p at which they expect to reach the end position. Sometimes, this is the same as the max end time T_a^{max} by when the agent is *required* to be at the end positions, but sometimes the preferred end time T_a^p can be sooner. Deviation from the preferred end time is modelled as a quadratic cost

$$C_a^T(s_a, T_a^p, K_a^T) = K_a^T \frac{1}{2} (t_a^n - T_a^p)^2 \quad (6)$$

incentivizing the agents to arrive at their preferred end time T_a^p by keeping the actual end time for each agent, t_a^n , close to T_a^p .

3.3.4. Regularizing cost

We find that adding a regularizing term to penalize extremely short time segments improves the overall quality of the paths by reducing near instantaneous motions in time. To penalize very fast agent motion, we use the regularizing term

$$C_a^R(s_a, K_a^R) = K_a^R \sum_{i=0}^n \left(\frac{t_a^n}{t_a^{i+1} - t_a^i} \right). \quad (7)$$

This is important because it allows us to feed the solver an initial space-time curve such as the ones in Figure 11 with many near instantaneous agent motions, and the solver is able to optimize the final space-time curve to a much more reasonable trajectory.

3.4. Interaction cost

3.4.1. Collision interaction cost

Interactions include agent-environment interactions and agent-agent interactions for which we must introduce a new arbitrary agent indexed by b where $a \neq b$. All our interactions depend on each




Figure 7: Tunneling artefact.




Figure 8: Resolving intersections between agents space-time rods resolves collisions in the scene since no agent shares the same x, y, t location at any point in the scene.

agent's 'radius of comfort' (or collision radius), denoted by r_a, r_b for agents a and b . The collision radius might be based on size, or a combination of size and personality: for example, people stay away from angry people. We encode these agent characteristics into our cost functions by extruding circle with radius r_a along the agent's path s_a thus forming a rod. In a scene with multiple agents, as long as no two rods are intersecting, the scene is collision-free as shown in Figure 8. The collision radius ensures that no two agents are in the same place at the same time. So even though the forces between the two space-time rods might be symmetric, the agents' response to these forces (change in path) leads to heterogeneous interactions. For static object interactions, the environment boundary is encoded into boundary vertices b_v and boundary elements b_e . Let us examine how we detect which rods are intersecting and how we deal with our three interaction types: agent-agent collisions, agent-agent gatherings (friendships) and agent-environment collisions.

Given a pair of agents a, b with 3D space-time curves s_a, s_b and collision radii r_a, r_b , we densely sample (upsample) each space-time curve uniformly. Next, we calculate the minimum smooth distances between the upsampled centerlines using the method described in Section 3.5 Equation (11). As long as the smooth minimum distance $dist$ between the upsampled rods $u_a = \text{upsample}(s_a), u_b = \text{upsample}(s_b)$ is further apart than $r_a + r_b$, collisions will not occur. We implement this non-linear constraint on




Figure 9: The log barrier energy quickly tends towards infinity as the minimum distance approaches the collision radius.

the agent paths as a log barrier energy:

$$C_{a,b}^C(\alpha, \mathbf{s}_a, \mathbf{s}_b, r_a, r_b, K_{a,b}^C) = -K_{a,b}^C \log(-r_a + r_b + \text{dist}(\alpha, \text{upsample}(\mathbf{s}_a), \text{upsample}(\mathbf{s}_b))). \quad (8)$$

where $K_{a,b}^C$ is the pair-wise weighting coefficient on the energy. The intuition behind a log barrier energy is explained in Figure 9.

Upsampling the trajectories prevents tunnelling artefacts (Figure 7) between edges during collision detection. The collision cost increases exponentially as the minimum distance between the two agent rods approaches the sum of the collision radii, so collisions are exponentially penalized. By summing this cost with the kinetic energy term, which incentivizes agents to move at a constant velocity, we get smooth local collision resolution as shown in Figure 8.

We employ two methods for sparsifying interactions. First, we use a 3D tree structure to store agent trajectory nodes (analogous to a Bounded Volume Hierarchy [BVH]) for broad-phase collision detection between trajectories. The tree-based broad phase reduces collision detection costs from $O(n^2)$ to $O(n \log(n))$. Second, if two agents are known to be far apart, we entirely avoid the collision detection step between those two agents thus reducing costs to $O(n)$ in the best case. Using a combination of tree-based broad phase along with manual denotation of interacting agents, we find that our method scales nearly as $O(n)$ as the number of agents in the scene increases.

3.4.2. Gathering interaction cost

While the collision resolution term pushes agent paths apart, we introduce a gathering term which pulls paths together

$$C_{a,b}^G(\alpha, \mathbf{s}_a, \mathbf{s}_b, r_a, r_b, K_{a,b}^G) = K_{a,b}^G (\text{dist}(\alpha, \text{upsample}(\mathbf{s}_a), \text{upsample}(\mathbf{s}_b)) - (r_a + r_b))^2. \quad (9)$$

This quadratic gathering term allows us to model scenarios in which friends going in the same direction tend to stick together as shown in Figure 18, or scenarios in which an agent needs to deliver a message

to another agent who is out of the way from his final destination such as Figure 20. The weighting term $K_{a,b}^G$ controls the desire of the agents a, b to group together.

3.4.3. Agent–environment interaction

We surpass previous methods in three ways in terms of agent–environment interactions. First, our method works on agents traversing highly constrained environments such as Figure 15a, 15b, or 19a. Second, our method allows agents to explore routes when multiple paths are available and choose the most optimal path using a novel modification to Dijkstra’s algorithm. Third, our path planning step is topology aware, so this method will work on more fascinating terrain.

Environment maps are stored as 2D mesh files with vertices and faces. We pre-compute the vertices and edges around the boundaries of the map and any static obstacles into b_v and b_e as a one-time preprocessing step. Next, we compute the minimum smooth distance between the map boundary edges and the agent rod. Luckily since static obstacles are fixed in space for all time, we can ignore the third dimension and only compute the 2D minimum distance, $d(\alpha, \mathbf{s}_a, b_v)$ from the rod’s projection onto the map and the map’s boundary edges. Next we pass the smooth distance d into the log barrier function

$$C_a^M(\alpha, \mathbf{s}_a, b_v, r_a) = -K_a^M \log(-r_a + \text{dist}(\alpha, \text{upsample}(\mathbf{s}_a), b_v)) \quad (10)$$

where K_a^M is the energy coefficient. Like in the agent–agent collision function (Equation 8), we can use the collision radius r_a for each agent since the initial agent paths are sufficiently far enough away from static obstacles such that $d > r_a$. The coarseness of the map mesh does not impact the smoothness of agent path; however, Dijkstra’s algorithm finds the shortest distance between two vertices on the map which correspond to the agent’s start and end positions to generate the initial path. Therefore, we set the nearest vertices to the given start and end points as the boundary conditions for the solve.

3.5. Smooth min distance implementation

The absolute minimum distance between agent paths is not a differentiable function. Therefore, we use a differentiable LogSumExp smooth minimum distance function to smoothly approximate the distance between two upsampled space–time curve vertices u_a and u_b . See the supplementary material for derivatives of the function. The smooth minimum distance function is

$$\text{dist}(\alpha, u_a, u_b) = \frac{-1}{\alpha} \log \left(\sum_{i=0}^n \sum_{j=0}^n e^{-\alpha \|s_a^i - s_b^j\|} \right) \quad (11)$$

where the α term controls the numerical sensitivity of the distance which changes depending on the size of the environment mesh and the distances between the agents’ path curves. A higher α leads to a more accurate minimum distance, but reduces numerical stability. Since this smooth minimum distance function is an underestimation of the true minimum distance, it is possible to get negative distances for very close objects.

Algorithm 1. A heuristic to find a usable α to be used within the agent-collision and map interaction cost, gradient and hessian functions.

```

function ALPHAHEURISTIC( $\alpha_0, s_a, s_b$ )
   $\alpha \leftarrow \alpha_0$ 
   $D \leftarrow \text{dist}(\alpha, s_a, s_b)$ 
  while  $D \leq 0$  do
     $\alpha \leftarrow \alpha + 0.1 * \alpha_0$ 
     $D \leftarrow \text{dist}(\alpha, s_a, s_b)$ 
  end while
  return  $\alpha$ 
end Function

```

We use Algorithm 1 to update α during the solve to guarantee a usable (real, non-negative) smooth minimum distance. Our initial α_0 for the scene is as low as possible, but large enough to guarantee that both our log barrier interaction costs are real and defined. If the initial α_0 is too small, the distance underestimation is too great, and the log barrier costs are either complex or undefined, then the interior point solver will throw an error. We must choose a large enough α_0 that the heuristic will provide an α that guarantees a real smooth distance values where $D > r_a + r_b$ to ensure real log barrier costs. For subsequent iterations, any real, non-negative smooth minimum distance is fine.

3.6. Initial paths

With our interior point solver, we require feasible initial trajectories in which the agents do not collide with other agents or any obstacles. The supplementary material shows our failed experiments in generating initial paths based on shortest distances. Some global path planner is necessary to generate feasible initial paths. Initially intersecting space-time rods causes NaNs in the log barrier interaction costs which makes the interior point method intractable. In order to generate initial plausible paths to feed into the optimization, we sequentially run a modified Dijkstra's algorithm for all agents traversing the environment mesh inspired by Treuille *et al.* [TCP06]. One might choose an A* heuristic to set initial paths, but we find that the cost of using an exact algorithm, like Dijkstra's, produces better results without the need to justify any heuristic. Dijkstra's is more costly than an inexact metric, but cost of finding initial paths is a very small part of the overall optimization.

Dijkstra finds the weighted least costly path along the edges of the map mesh from the start location to the end location for a given agent. The edge weights are set based on spatial distance and time away from T^{\max} to incentivize that agents reach the final spatial location in as little time as possible in order to avoid blocking the paths of other agents. We avoid obstacles (environment boundary) by giving these edges, b_e , a prohibitively high traversal cost. We only consider map edges that have a distance from the map boundary at least greater than the agent's collision radius.

Furthermore, we ensure that every agent has a collision-free initial path by making it impossible for an agent's initial path to pass too close to the initial path of another agent. For each agent, upon finding an initial path, we wipe out all edges connected to the path within the agent's collision radius. The process of sequen-




Figure 10: Running Dijkstra's algorithm to set initial agent paths on the map. Red map edges indicate a higher traversal cost. Since we want to set a feasible initial path for each agent, we keep them away from the map edges as well as away from other agent's initial paths.

tially generating initial paths for each agent impacts the final results; therefore, agent order cannot be entirely ignored. Changing the order could change the optimization as well. For very constrained scenes, it becomes impossible to ensure collision-free initial trajectories since agents often run out of traversible edges from $[x^0, y^0]$ to $[x^n, y^n]$. Since agents move through both space and time, we extrude our environment map into the t dimension and create a pre-set number of layers from t^0 to T^{\max} as shown in Figure 10. Solving this 3D space-time Dijkstra's algorithm makes it a lot more feasible to find collision-free initial paths for all agents. Sometimes, the coarseness of the environment mesh, or the low number of layers in the 3D space-time map makes it impossible for Dijkstra's algorithm to find a feasible initial path. In this case, either the environment map would need to be sub-divided, more layers would need to be added, or the scene itself might be infeasible given the T^{\max} time constraints and the number and size of agents involved. Additionally, although the space-time Dijkstra's algorithm will provide feasible initial paths, they will have many kinks and sharp turns (Figure 11) indicating unnatural motion and thus cannot be used in the simulation directly. These get smoothed out during the optimization process resulting in much more realistic motion. It is possible that this proposed method to generate initial paths will affect the state of the final trajectories. For example, if an initial path passes through the left of an obstacle, it might be stuck in a local minimum even though an optimal trajectory might be through the right side. This drawback is somewhat mitigated by using Dijkstra's search to pre-compute the least costly paths. Either way, whether or not the solver produces a globally minimum outcome, the trajectories are guaranteed to be smooth, viable and collision free.

3.7. Optimization

Now aggregating over all the agents in the scene, $\|A\|$, we put intra-agent cost functions equations (4)– (7) together with interaction costs equations (8)– (10) and construct an optimization problem for our scene. We aggregate boundary conditions for each agent and




Figure 11: Top is the initial space-time curve output from Dijkstra's path finding algorithm. Bottom is the final result of the solver with optimal paths. Agents are allowed to find their own optimal end-times while avoiding collisions in tightly constrained bottle-necks such as this.

denote them with capital letters to show that they apply to the entire scene. The optimization problem

$$f^* = \min \sum_{a=1}^{\|A\|} C_a^K + C_a^A + C_a^T + C_a^R + C_a^M + \quad (12a)$$

$$\sum_{\substack{a,b \in \\ \{1..|A|\} \\ |a \neq b}} C_{a,b}^G + C_{a,b}^C \quad (12b)$$

$$s.t. T^i \leq T^{i+1} \quad (12c)$$

$$(X^0, Y^0, T^0) = (X, Y, T)^{start} \quad (12d)$$

$$(X^n, Y^n) = (X, Y)^{end} \quad (12e)$$

$$T^n \leq T^{max} \quad (12f)$$

is solved using a quasi-Newton interior point method. An outer loop over this optimization is required for a proper log-barrier method [NW06]. The full pseudocode overview of our method's outer loop, 2, illustrates this. The supplementary material for our method include the sub-routines, the calculations for the gradients and Hessian approximations of our costs, intuition for controlling agent behaviour through function weights.

3.8. Controlling agent behaviour

Controlling agent behaviour is very intuitive in our method. There are several points of control in our simulation. First, each agent has its own pre-determined characteristics profile. This includes things such as size, mass (doubles as stubbornness), collision radius, grouping preferences, preferred arrival time. Each of these parameters can be intuitively adjusted for any agent to get the desired be-

Algorithm 2. A full pseudocode overview of our method with sub-routines provided in the supplementary material.

global input variables

$OuterIts \geq 1$, outer loop iterations
 $\mu, c = 0.75$, log-barrier coeff and its decrement factor
 $\alpha_0, cutoff = 0.2$, initial alpha, Hess sparsifying cutoff
 T^{max} , max time constraint for agents
 b_e, b_v , environment edges and vertices.
 K , cost function weights for agents
 R , collision radius for agents
 T , preferred end times
 M , agent masses

end global input variables

interior point solver parameters

$MaxIts$, max iterations
 B , trajectory boundary conditions
 Stop if $StepSize > 10^{-2}$ (metres), norm of solver step
 Stop if $FirstOrderOpt > 10^{-2}$ (metres), first-order optimality criteria

end interior point solver parameters

Require: @COSTS, @GRADS, @HESS (supplementary material)

function FINDOPTIMALPATHS

```
[S, B] ← DJIKSTRASPREPROCESS(R, T, b_v)
S ← increment t by ε to ensure t_{i+1} > t_i
while OuterIts > 0 do
  S ← IPSOLVER(S, B, @COSTS, @GRADS, @HESS)
  μ ← cμ
  OuterIts ← OuterIts - 1
end while
return S
end Function
```

haviour as shown in Figure 4. Another control point for this method is adjusting the coefficients of the energies as shown in Figure 13. If the scene happens to be on a large, unconstrained map, its advantageous to turn $K^M = 0$ off to prevent possible numerical issues from the smooth distance function. Agents are allowed to collide if K^C is set to 0. The gathering cost can be modified through the grouping parameter for each agent pair K^G . If agents need to reach their goal by a specific end time, we use K^T to control this behaviour. Acceleration is controlled by the K_A parameter. Regularizing the edge lengths of each curve is important, as explained above, but modifying K^R can affect the likelihood of the solver getting stuck in local minimums in highly constrained environments. Last, but not least, K^K weights the kinetic energy term and increasing it (or increasing the mass term m) will incentivize the agent to follow a more direct path to the end location.

4. Results

The performance of our method relies on (1) the number of agents in the scenes and (2) the complexity of the agent's paths through space and time. We push along each of these axes separately in this section. We show extremely complex maze-like environments with bottlenecks, as well as large scenes with tens (to hundreds) of agents. In addition to the comparisons shown in the related works, we include the obligatory circles of agents in Figure 14. Notice that our agents take smooth, natural trajectories rather than the strange spinning motions demonstrated by other methods. Next, in Figure 18, we




Figure 12: (Left) Varying number of DOFs per agent from 90 to 540 on a fixed scene with eight agents in a circle (like Figure 14). Every agent interacts with every other agent in the scene. (Middle) Varying number of agents with 300 DOFs per agent in a circle (like Figure 14). Every agent interacts with every other agent in the scene. (Right) Varying DOFs in the environment boundary from 2520 to 20,160 on the corn maze mesh with three agents of 600 DOFs each in the scene. Scaling tests show **nearly** linear performance in the worst-case shown in these plots and linear performance in the best case.




Figure 13: Updating the energy coefficients to reach a desirable scene is a simple and intuitive process in our method. This figure shows the simplicity of the process, and illustrates that a desirable scene can be reached via many different parameter configurations. In the top row, increasing collision costs reduces the likelihood of collision between agents. In the middle row, increasing kinetic costs incentivizes more direct paths. In the bottom row, increasing acceleration costs locally decreases lateral motion.

highlight the several different types of heterogeneous interactions supported by our method. We show the naive case where agents collide. We show standard symmetric collision between agents. We show a stubborn snake which forces the scared humans to move out of its path. We show a large elk who is still scared of the humans so it changes its trajectory just as the human agents change theirs. We show a large and stubborn elephant that goes straight to its destination while the humans have to significantly alter trajectories to avoid it. Finally, we show a scenario where two friends stick closer to each other on their way to their final destination.

4.1. Tight spaces

In Figures 15 and 16, we show our method works in extremely tight environments. ORCA and Implicit Crowds show that they can navi-

gate environmental constraints with manual pre-processing to define trajectories around obstacles so agents do not bump into them, but nowhere near as tight as our examples. Meanwhile, NHTTC does not implement environmental constraints altogether. We show that our method works in extremely constrained scenarios such as a subway tunnel or a tight warehouse of robots where each lane is big enough to accommodate only a single agent and agents have to take turns to pass through to avoid locking.

4.2. Flexible arrivals

Although our method is built to handle heterogeneous interactions and agents in highly constrained spaces, we also show several other features of our method. First, our agents have flexible arrival times. In Figure 11, we show a bottleneck where all the agents simply cannot arrive at their destinations by their preferred end times. This feature serves as a counterpoint to the idea of using fixed start and end boundary conditions to model trajectory curves. In Figure 16, we show a highly constrained bottleneck of agents trying to get to their seats on an airplane.

4.3. Complex environments

Additionally, we show the navigability of agents in larger constrained environments in Figures 19a and 19b. Agents are able to navigate the mazes while avoiding collisions with each other at the bottlenecks.

4.4. Controllability

In Figure 20, we show the flexibility and controllability of our method. We create a scenario where one agent must meet up with another agent to deliver a message and then arrive at his end goal faster than the other agents. In another scenario, we direct the agent to meet up with the second agent, yet arrive at his end goal at the same time as the other agents. Last, we also provide the simple scene where agents are ignorant of each other. The flash mob example in Figure 21 shows how our method can be used to position agents intricately. The corresponding submission video shows how we are able to control the order of the placement of the agents in the flash mob as well through the preferred end time parameter. Our method





Figure 14: *Circles of agents.*



(a) Maze of roombas.

(b) Subway platform

Figure 15: *In both (a) and (b) only one agent can move through the intersection at a time.*




Figure 16: *Chaos in first class. All the passengers are in incorrect seats.*




Figure 17: *A large scene with 102 agents moving to opposite sides of a meadow while a couple of deer avoid them.*

allows careful control with no manual effort on the part of the user, thus making it useful for games or other industrial applications.

4.5. Larger scenes

And finally, for the sake of completeness, we show Figure 17, a large scene with 104 agents. We include this example to demonstrate that even though our method is designed for small-to-medium-sized scenes, it can work on simple scenes with larger numbers of agents with heterogeneous interactions and still scales linearly in time.

4.6. Timings

Figure 12 provides three plots, each with different scaling information for different usage scenarios. The first plot measure the average time per solver function iteration for an increasing number of agents. Each agent trajectory has 100 nodes (300 DOFs) and the performance is near linear with broad-phase collision detection enabled. Our second plot measure the average time per iteration for an increasing number of DOFS for an eight-agent circle. The number of nodes starts off at 30 and goes all the way to 180 nodes per agent. Performance is again nearly linear with broad phase collision detection enabled. Without a broad phase, performance is quadratic. Something to note is that increasing the number of DOFs per agent provides no additional benefit to the quality of the simulation since the actual physics of the space-time rods are not the end result of our method. As long as agents trajectory curves have enough DOFs to traverse the environment smoothly, the end results will be good. Our third plot measures the time per iteration for an increasing complexity in the environment mesh on the corn maze (Figure 19a) example with three agents and 200 trajectory nodes per agent. Again, performances are nearly linear.

As mentioned before, there is no gold-standard crowd simulation algorithm since each scenario is so unique and intricate. Rather than focus on timing performance for large crowds, our method focuses on solving previously overlooked path planning scenarios with heterogeneous interactions for small-to-medium-sized groups in highly constrained environments. Our asymptotics show linear to near-linear performance, but there is plenty of room for improvement in wall-clock-times through optimization and parallelization.

5. Conclusion and Future Work

In this paper, we show that modelling the motion of agents through space and time using a 3D curve resolves a number of difficulties with multi-agent path planning. Assigning physical characteristics such as a radius and mass to the agent's trajectory curve lets us in-





Figure 18: (0) No interactions. (1) Mass (analogous to stubbornness) weighted heterogeneous interactions. (2) Size and mass-weighted interactions. (3) Symmetric interactions.(4) Friends rendezvous along the way. (5) Size-based interactions.



(a) Corn maze

(b) Circle maze

Figure 19: (a) Large cornmaze with eight agents all trying to get to different locations through the maze. (b) A circular-shaped maze where agents need to navigate while avoiding collisions.





Figure 20: Left: An agent delivers a message to another agent and rushes to his goal position. Middle: Two agents meet up briefly before walking to their separate destinations. Right: No contact between any agents.

tuitively simulate heterogeneous interactions between agents. Our method outputs agent paths that are smooth, can navigate through highly constrained environments and are parameterized by intuitive controls.

In the future, we hope to improve our agent model to include limited environmental perception (rather than the omniscience our agents enjoy currently) as well as additional dynamics. As men-



(a) Pre-smile flash mob

(b) Post-smile

Figure 21: (a) Group of people in a flash mob instructed to form a smily face. Agents follow unintuitive paths through space and time to maintain the fluidity of their motion. (b) A smily face created by controlling the motion of the mob.

tioned before, performance depends on two factors: number of agents and their temporal support. A time resolution that is too sparse will not allow Dijkstra's to produce feasible initial trajectories. Increasing the time resolution overcomes this limitation. While we can currently push along these axes independently, a future work is to be able to push along both axes together, *i.e.* handle complex environments with large numbers of agents with more intricate and more realistic motion. We also hope to extend the method in order to handle scenes with conflicting and changing goals, *e.g.* a predator-prey scenario.

We would also like to extend our space-time approach to fully 3D environments which requires performing our optimization in four dimensions. Finally, while our focus is on robustness in path planning, there is still further room for improvement in performance. We are excited to explore fast space-time multi-agent path planning to scale our approach to the large dense crowds that continuum approaches excel at, while maintaining our unique advantages.

Acknowledgements

This work is funded in part by the National Science Foundation, Connaught Fund, CFI-JELF Fund Accelerator, the Ontario Early Research Award program and the Canada Research Chair Program as well as gifts by Adobe Systems. Additionally, we thank Sarah Kushner for lending us her voice for the video and for proofreading. We also thank members of the DGP lab for participating in the live motion-capture scenes included in our results. Finally, we thank our anonymous reviewers for helpful feedback.

References

- [Aum87] AUMANN R. J.: Correlated equilibrium as an expression of Bayesian rationality. *Econometrica: Journal of the Econometric Society* 55 (1987), 1–18.
- [CC14] CHARALAMBOUS P., CHRYSANTHOU Y.: The PAG crowd: A graph based approach for efficient data-driven crowd simulation. *Computer Graphics Forum* 33 (2014), 95–108.
- [CvTZ*22] COLAS A., VAN TOLL W., ZIBREK K., HOYET L., OLIVIER A., PETTRÉ J.: Interaction fields: Intuitive sketch-based steering behaviors for crowd simulation. *Computer Graphics Forum*, 41, 2 (2022), 521–534.
- [Dav20] DAVIS B.: NHTTC source code. <https://github.com/davisbo/NHTTC> (2020). Accessed Jan 10, 2022.
- [DKG20] DAVIS B., KARAMOUZAS I., GUY S.: NH-TTC: A gradient-based framework for generalized anticipatory collision avoidance. In *Robotics: Science and Systems* (2020). <https://doi.org/10.15607/RSS.2020.XVI.078>
- [DMCN*17] DUTRA T. B., MARQUES R., CAVALCANTE-NETO J. B., VIDAL C. A., PETTRÉ J.: Gradient-based steering for vision-based crowd simulation algorithms. *Computer Graphics Forum* 36 (2017), 337–348.
- [FR13] FLAGG M., REHG J. M.: Video-based crowd synthesis. *IEEE Transactions on Visualization and Computer Graphics* 19, 11 (2013), 1935–1947. <https://doi.org/10.1109/TVCG.2012.317>
- [GCC*10] GUY S. J., CHHUGANI J., CURTIS S., DUBEY P., LIN M. C., MANOCHA D.: PLEdetrans: A least-effort approach to crowd simulation. In *Proceedings of the 2010 Eurographics/ACM SIGGRAPH Symposium on Computer Animation, SCA* (Madrid, Spain, 2010), Z. Popovic and M. A. Otaduy (Eds.), Eurographics Association, pp. 119–128. <http://doi.org/10.2312/SCA/SCA10/119-128>
- [GKLM11] GUY S. J., KIM S., LIN M. C., MANOCHA D.: Simulating heterogeneous crowd behaviors using personality trait theory. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2011), pp. 43–52.
- [GNCL14] GOLAS A., NARAIN R., CURTIS S., LIN M. C.: Hybrid long-range collision avoidance for crowd simulation. *IEEE Transactions on Visualization and Computer Graphics* 20, 7 (2014), 1022–1034. <https://doi.org/10.1109/TVCG.2013.235>
- [GVDBL*12] GUY S. J., VAN DEN BERG J., LIU W., LAU R., LIN M. C., MANOCHA D.: A statistical similarity measure for aggregate crowd dynamics. *ACM Transactions on Graphics (TOG)*, 31, 6 (2012), 1–11.
- [HKHL13] HYUN K., KIM M., HWANG Y., LEE J.: Tiling motion patches>0. *IEEE Transactions on Visualization and Computer Graphics* 19, 11 (2013), 1923–1934. <https://doi.org/10.1109/TVCG.2013.80>
- [HM98] HELBING D., MOLNAR P.: Social force model for pedestrian dynamics. *Physical Review E* 51, (May 1998), 4282–4286. <https://doi.org/10.1103/PhysRevE.51.4282>
- [Hug02] HUGHES R.: A continuum theory for the flow of pedestrians. *Transportation Research Part B: Methodological* 36, (July 2002), 507–535. [https://doi.org/10.1016/S0191-2615\(01\)00015-7](https://doi.org/10.1016/S0191-2615(01)00015-7)
- [KGvdS13] KARAMOUZAS I., GERAERTS R., van der STAPPEN A. F. (2013) Space-time group motion planning. In: Frazzoli E., Lozano-Perez T., Roy N., Rus D. (Eds.) *Algorithmic Foundations of Robotics X*. Springer Tracts in Advanced Robotics, vol 86. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-36279-8_14
- [KSNG17] KARAMOUZAS I., SOHRE N., NARAIN R., GUY S. J.: Implicit crowds: Optimization integrator for robust crowd simulation. *ACM Transactions on Graphics*, 36, 4 (July 2017). <https://doi.org/10.1145/3072959.3073705>
- [LLKP11] LEVINE S., LEE Y., KOLTUN V., POPOVIC Z.: Space-time planning with parameterized locomotion controllers. *ACM Transactions on Graphics* 30, (May 2011). <https://doi.org/10.1145/1966394.1966402>
- [NGCL09] NARAIN R., GOLAS A., CURTIS S., LIN M. C.: Aggregate dynamics for dense crowd simulation. In *ACM SIGGRAPH Asia 2009 Papers* (2009), pp. 1–8.
- [NW06] NOCEDAL J., WRIGHT S. J.: *Numerical Optimization* (2nd edition). Springer, New York, NY, USA, 2006.
- [NX15a] Neo-X: PPRParameters. 2015 <https://github.com/SteerSuite/Release/blob/master/pprAI/include/PPRParameters.h> (2015). Accessed Oct 30, 2022.
- [NX15b] Neo-X: RVO2D parameters. https://github.com/SteerSuite/Release/blob/master/rvo2AI/include/RVO2D_Parameters.h (2015). Accessed Oct 30, 2022.
- [NX15c] Neo-X: Socialforces parameters. https://github.com/SteerSuite/Release/blob/master/socialForcesAI/include/SocialForces_Parameters.h (2015). Accessed Oct 30, 2022.
- [OPOD10] ONDŘEJ J., PETTRÉ J., OLIVIER A.-H., DONIKIAN S.: A synthetic-vision based steering approach for crowd simulation. *ACM Transactions on Graphics (TOG)*, 29, 4 (2010), 1–9.
- [PGOSB05] PELECHANO GÓMEZ N., O'BRIEN K., SILVERMAN B. G., BADLER N.: Crowd simulation incorporating agent

- psychological models, roles and communication. In *Proceedings of the First International Workshop on Crowd Simulation* (2005).
- [PSE*00] POPOVIĆ J., SEITZ S. M., ERDMANN M., POPOVIĆ Z., WITKIN A.: Interactive manipulation of rigid body simulations. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (2000), pp. 209–217.
- [RGHD13] RAMOS F., GAJAMOHAN M., HUEBEL N., D'ANDREA R.: Time-optimal online trajectory generator for robotic manipulators (2013).
- [RLC*14] RAMIREZ J. G. R., LANGE D., CHARALAMBOUS P., ESTEVES C., PETTRÉ J.: Optimization-based computation of locomotion trajectories for crowd patches. In *MIG'14: Proceedings of the Seventh International Conference on Motion in Games* (New York, NY, USA, 2014), Association for Computing Machinery, pp. 7–16. <https://doi.org/10.1145/2668064.2668094>
- [SKF08] SHAPIRO A., KALLMANN M., FALOUTSOS P.: Interactive motion correction and object manipulation. In *ACM SIGGRAPH 2008 Classes* (2008), pp. 1–8.
- [SKFR09] SINGH S., KAPADIA M., FALOUTSOS P., REINMAN G.: SteerBench: A benchmark suite for evaluating steering behaviors. *Computer Animation & Virtual Worlds*, 20, 5–6 (Sep. 2009), 533–548.
- [SKH*11] SINGH S., KAPADIA M., HEWLETT B., REINMAN G., FALOUTSOS P.: A modular framework for adaptive agent-based steering. In *Proceedings of the Symposium on Interactive 3D Graphics and Games* (2011), pp. 141–150.
- [Sta14] STADLER J. J.: A Framework for Video-driven Crowd Synthesis. Thesis, 2014.
- [TCP06] TREUILLE A., COOPER S., POPOVIĆ Z.: Continuum crowds. *ACM Transactions on Graphics (TOG)*, 25, 3 (2006), 1160–1168.
- [TYK*09] TAKAHASHI S., YOSHIDA K., KWON T., LEE K. H., LEE J., SHIN S. Y.: Spectral-based group formation control. *Computer Graphics Forum* 28 (2009), 639–648.
- [Unk15] KARAMOUZAS I. Implicit. <https://github.com/johnoriginal/implicit-crowds/blob/master/data/implicit.ini> (2015). Accessed Jan 10, 2022.
- [VDBGML11] VAN DEN BERG J., GUY S. J., LIN M., MANOCHA D.: (2011) Reciprocal n-body collision avoidance. In *2011 Robotics Research*. Springer Tracts in Advanced Robotics, Springer, Berlin, Heidelberg, vol. 70, 31.
- [vdBO04] VAN DEN BERG J., OVERMARS M.: Roadmap-based motion planning in dynamic environments. In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)* (2004), vol. 2, pp. 1598–1605. <https://doi.org/10.1109/IROS.2004.1389624>
- [vdBO05] VAN DEN BERG J., OVERMARS M.: Prioritized motion planning for multiple robots. In *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2005), pp. 430–435. <https://doi.org/10.1109/IROS.2005.1545306>
- [VTP19] VAN TOLL W., PETTRÉ J.: Connecting global and local agent navigation via topology. In *Proceedings of the 12th ACM SIGGRAPH Conference on Motion, Interaction and Games* (2019), pp. 33:1–33:10.
- [VTTK*16] VAN TOLL W., TRIESSCHEIJN R., KALLMANN M., OLIVA R., PELECHANO N., PETTRÉ J., GERAERTS R.: A comparative study of navigation meshes. In *Proceedings of the 9th International ACM SIGGRAPH Conference on Motion in Games* (2016), pp. 91–100.
- [WBHB18] WINKLER A. W., BELLICOSO D. C., HUTTER M., BUCHLI J.: Gait and trajectory optimization for legged systems through phase-based end-effector parameterization. *IEEE Robotics and Automation Letters (RA-L)* 3 (July 2018), 1560–1567. <https://doi.org/10.1109/LRA.2018.2798285>
- [WJGO*14] WOLINSKI D. J., GUY S., OLIVIER A.-H., LIN M., MANOCHA D., PETTRÉ J.: Parameter estimation and comparative evaluation of crowd simulations. *Computer Graphics Forum* 33, (2014), pp. 303–312.
- [WK88] WITKIN A., KASS M.: Spacetime constraints. *ACM SIGGRAPH Computer Graphics*, 22, 4 (1988), 159–168.
- [WLJT17] WEISS T., LITTENEKER A., JIANG C., TERZOPOULOS D.: Position-based multi-agent dynamics for real-time crowd simulation. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2017), pp. 1–2.
- [WLP16] WOLINSKI D., LIN M. C., PETTRÉ J.: WarpDriver: Context-aware probabilistic motion prediction for crowd simulation. *ACM Transactions on Graphics (TOG)*, 35, 6 (2016), 1–11.
- [YL16] YU J., LAVALLE S. M.: Optimal multirobot path planning on graphs: Complete algorithms and effective heuristics. *IEEE Transactions on Robotics*, 32, 5 (2016), 1163–1177.
- [YMPT09] YERSIN B., MAIM J., PETTRÉ J., THALMANN D.: Crowd patches: Populating large-scale virtual environments for real-time applications. In *3D'09: Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2009), Association for Computing Machinery, pp. 207–214. <https://doi.org/10.1145/1507149.1507184>
- [YSC21] YU C., SCHUMACHER H., CRANE K.: Repulsive curves. *ACM Transactions on Graphics* 40, 2 (2021), 1–21.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Supporting Information