
A message-passing algorithm for multi-agent trajectory planning

José Bento*

jbento@disneyresearch.com

Nate Derbinsky

nate.derbinsky@disneyresearch.com

Javier Alonso-Mora

jalonso@disneyresearch.com

Jonathan Yedidia

yedidia@disneyresearch.com

Abstract

We describe a novel approach for computing collision-free *global* trajectories for p agents with specified initial and final configurations, based on an improved version of the alternating direction method of multipliers (ADMM). Compared with existing methods, our approach is naturally parallelizable and allows for incorporating different cost functionals with only minor adjustments. We apply our method to classical challenging instances and observe that its computational requirements scale well with p for several cost functionals. We also show that a specialization of our algorithm can be used for *local* motion planning by solving the problem of joint optimization in velocity space.

1 Introduction

Robot navigation relies on at least three sub-tasks: localization, mapping, and motion planning. The latter can be described as an optimization problem: compute the lowest-cost path, or *trajectory*, between an initial and final configuration. This paper focuses on trajectory planning for *multiple* agents, an important problem in robotics [1, 2], computer animation, and crowd simulation [3].

Centralized planning for multiple agents is PSPACE hard [4, 5]. To contend with this complexity, traditional multi-agent planning prioritizes agents and computes their trajectories sequentially [6], leading to suboptimal solutions. By contrast, our method plans for all agents simultaneously. Trajectory planning is also simplified if agents are non-distinct and can be dynamically assigned to a set of goal positions [1]. We consider the harder problem where robots have a unique identity and their goal positions are statically pre-specified. Both mixed-integer quadratic programming (MIQP) [7] and [more efficient, although local] sequential convex programming [8] approaches have been applied to the problem of computing collision-free trajectories for multiple agents with pre-specified goal positions; however, due to the non-convexity of the problem, these approaches, especially the former, do not scale well with the number of agents. Alternatively, trajectories may be found by sampling in their joint configuration space [9]. This approach is probabilistic and, alone, only gives asymptotic guarantees. See Appendix A for further comments on discrete search methods.

Due to the complexity of planning collision-free trajectories, *real-time* robot navigation is commonly decoupled into a global planner and a fast local planner that performs collision-avoidance. Many single-agent reactive collision-avoidance algorithms are based either on potential fields [10], which typically ignore the velocity of other agents, or “velocity obstacles” [11], which provide improved performance in dynamic environments by formulating the optimization in velocity space instead of Cartesian space. Building on an extension of the velocity-obstacles approach, recent work on centralized collision avoidance [12] computes collision-free local motions for all agents whilst maximizing a joint utility using either a computationally expensive MIQP or an efficient, though local, QP. While not the main focus of this paper, we show that a specialization of our approach

*This author would like to thank Emily Hupf and Noa Ghersin for their support while writing this paper.

to global-trajectory optimization also applies for local-trajectory optimization, and our numerical results demonstrate improvements in both efficiency and scaling performance.

In this paper we formalize the global trajectory planning task as follows. Given p agents of different radii $\{r_i\}_{i=1}^p$ with given desired initial and final positions, $\{x_i(0)\}_{i=1}^p$ and $\{x_i(T)\}_{i=1}^p$, along with a cost functional over trajectories, compute collision-free trajectories for all agents that minimize the cost functional. That is, find a set of intermediate points $\{x_i(t)\}_{i=1}^p$, $t \in (0, T)$, that satisfies the “hard” collision-free constraints that $\|x_i(t) - x_j(t)\| > r_i + r_j$, for all i, j and t , and that insofar as possible, minimizes the cost functional.

The method we propose searches for a solution within the space of piece-wise linear trajectories, wherein the trajectory of an agent is completely specified by a set of positions at a fixed set of time instants $\{t_s\}_{s=0}^\eta$. We call these time instants *break-points* and they are the same for all agents, which greatly simplifies the mathematics of our method. All other intermediate points of the trajectories are computed by assuming that each agent moves with constant velocity in between break-points: if t_1 and $t_2 > t_1$ are consecutive break-points, then $x_i(t) = \frac{1}{t_2 - t_1}((t_2 - t)x_i(t_1) + (t - t_1)x_i(t_2))$ for $t \in [t_1, t_2]$. Along with the set of initial and final configurations, the number of interior break-points $(\eta - 1)$ is an input to our method, with a corresponding tradeoff: increasing η yields trajectories that are more flexible and smooth, with possibly higher quality; but increasing η enlarges the problem, leading to potentially increased computation.

The main contributions of this paper are as follows:

- i) We formulate the global trajectory planning task as a decomposable optimization problem. We show how to solve the resulting sub-problems exactly and efficiently, despite their non-convexity, and how to coordinate their solutions using message-passing. Our method, based on the “three-weight” version of ADMM [13], is easily parallelized, does not require parameter tuning, and we present empirical evidence of good scalability with p .
- ii) Within our decomposable framework, we describe different sub-problems, called *minimizers*, each ensuring the trajectories satisfy a separate criterion. Our method is flexible and can consider different combinations of minimizers. A particularly crucial minimizer ensures there are no inter-agent collisions, but we also derive other minimizers that allow for finding trajectories with minimal total energy, avoiding static obstacles, or imposing dynamic constraints, such as maximum/minimum agent velocity.
- iii) We show that our method can specialize to perform local planning by solving the problem of joint optimization in velocity space [12].

Our work is among the few examples where the success of applying ADMM to find approximate solutions to a large non-convex problems can be judged with the naked eye, by the gracefulness of the trajectories found. This paper also reinforces the claim in [13] that small, yet important, modifications to ADMM can bring an order of magnitude increase in speed. We emphasize the importance of these modifications in our numerical experiments, where we compare the performance of our method using the three-weight algorithm (TWA) versus that of standard ADMM.

The rest of the paper is organized as follows. Section 2 provides background on ADMM and the TWA. Section 3 formulates the global-trajectory-planning task as an optimization problem and describes the separate blocks necessary to solve it (the mathematical details of solving these sub-problems are left to appendices). Section 4 evaluates the performance of our solution: its scalability with p , sensitivity to initial conditions, and the effect of different cost functionals. Section 5 explains how to implement a velocity-obstacle method using our method and compares its performance with prior work. Finally, Section 6 draws conclusions and suggests directions for future work.

2 Minimizers in the TWA

In this section we provide a short description of the TWA [13], and, in particular, the role of the *minimizer* building blocks that it needs to solve a particular optimization problem. Section B of the supplementary material includes a full description of the TWA.

As a small illustrative example of how the TWA is used to solve optimization problems, suppose we want to solve $\min_{x \in \mathbb{R}^3} f(x) = \min_{\{x_1, x_2, x_3\}} f_1(x_1, x_3) + f_2(x_1, x_2, x_3) + f_3(x_3)$, where $f_i(\cdot) \in$

$\mathbb{R} \cup \{+\infty\}$. The functions can represent soft costs, for example $f_3(x_3) = (x_3 - 1)^2$, or hard equality or inequality constraints, such as $f_1(x_1, x_3) = \mathbb{J}(x_1 \leq x_3)$, where we are using the notation $\mathbb{J}(\cdot) = 0$ if (\cdot) is true or $+\infty$ if (\cdot) is false.

The TWA solves this optimization problem iteratively by passing messages on a bipartite graph, in the form of a Forney factor graph [14]: one *minimizer-node* per function f_b , one *equality-node* per variable x_j and an edge (b, j) , connecting b and j , if f_b depends on x_j (see Figure 1-left).




Figure 1: Left: bipartite graph, with one minimizer-node on the left for each function making up the overall objective function, and one equality-node on the right per variable in the problem. Right: The input and output variables for each minimizer block.

Apart from the first-iteration message values, and two internal parameters¹ that we specify in Section 4, the algorithm is fully specified by the behavior of the minimizers and the topology of the graph.

What does a minimizer do? The minimizer-node g_1 , for example, solves a small optimization problem over its local variables x_1 and x_3 . Without going into the full detail presented in [13] and the supplementary material, the estimates $x_{1,1}$ and $x_{1,3}$ are then combined with running sums of the differences between the minimizer estimates and the equality-node consensus estimates to obtain *messages* $m_{1,1}$ and $m_{1,3}$ on each neighboring edge that are sent to the neighboring equality-nodes along with corresponding *certainty weights*, $\bar{\rho}_{1,2}$ and $\bar{\rho}_{1,3}$. All other minimizers act similarly.

The equality-nodes receive these local messages and weights and produce consensus estimates for all variables by computing an average of the incoming messages, weighted by the incoming certainty weights $\bar{\rho}$. From these consensus estimates, *correcting messages* are computed and communicated back to the minimizers to help them reach consensus. A certainty weight for the correcting messages, $\leftarrow{\rho}$, is also communicated back to the minimizers. For example, the minimizer g_1 receives correcting messages $n_{1,1}$ and $n_{1,3}$ with corresponding certainty weights $\leftarrow{\rho}_{1,1}$ and $\leftarrow{\rho}_{1,3}$ (see Figure 1-right).

When producing new local estimates, the b th minimizer node computes its local estimates $\{x_j\}$ by choosing a point that minimizes the sum of the local function f_b and weighted squared distance from the incoming messages (ties are broken randomly):

$$\{x_{b,j}\}_j = g_b(\{n_{b,j}\}_j, \{\leftarrow{\rho}_{b,j}^k\}_j) \equiv \arg \min_{\{x_j\}_j} \left[f_b(\{x_j\}_j) + \frac{1}{2} \sum_j \leftarrow{\rho}_{b,j} (x_j - n_{b,j})^2 \right], \quad (1)$$

where $\{ \}_j$ and \sum_j run over all equality-nodes connected to b . In the TWA, the certainty weights $\{\leftarrow{\rho}_{b,j}\}$ that this minimizer outputs must be 0 (*uncertain*); ∞ (*certain*); or ρ_0 , set to some fixed value. The logic for setting weights from minimizer-nodes depends on the problem; as we shall see, in trajectory planning problems, we only use 0 or ρ_0 weights. If we choose that all minimizers always output weights equal to ρ_0 , the TWA reduces to standard ADMM; however, 0-weights allows equality-nodes to ignore inactive constraints, traversing the search space much faster.

Finally, notice that all minimizers can operate simultaneously, and the same is true for the consensus calculation performed by each equality-node. The algorithm is thus easy to parallelize.

3 Global trajectory planning

We now turn to describing our decomposition of the global trajectory planning optimization problem in detail. We begin by defining the variables to be optimized in our optimization problem. In

¹These are the step-size and ρ_0 constants. See Section B in the supplementary material for more detail.

our formulation, we are *not* tracking the points of the trajectories by a continuous-time variable taking values in $[0, T]$. Rather, our variables are the positions $\{x_i(s)\}_{i \in [p]}$, where the trajectories are indexed by i and break-points are indexed by a discrete variable s taking values between 1 and $\eta - 1$. Note that $\{x_i(0)\}_{i \in [p]}$ and $\{x_i(\eta)\}_{i \in [p]}$ are the initial and final configuration, sets of fixed values, not variables to optimize.

3.1 Formulation as unconstrained optimization without static obstacles

In terms of these variables, the *non-collision* constraints² are

$$\begin{aligned} & \|(\alpha x_i(s+1) + (1-\alpha)x_i(s)) - (\alpha x_j(s+1) + (1-\alpha)x_j(s))\| \geq r_i + r_j, \\ & \text{for all } i, j \in [p], s \in \{0, \dots, \eta-1\} \text{ and } \alpha \in [0, 1]. \end{aligned} \quad (2)$$

The parameter α is used to trace out the constant-velocity trajectories of agents i and j between break-points $s+1$ and s . The parameter α has no units, it is a normalized time rather than an absolute time. If t_1 is the absolute time of the break-point with integer index s and t_2 is the absolute time of the break-point with integer index $s+1$ and t parametrizes the trajectories in absolute time then $\alpha = (t - t_1)/(t_2 - t_1)$. Note that in the above formulation, absolute time does not appear, and any solution is simply a set of paths that, when travelled by each agent at constant velocity between break-points, leads to no collisions. When converting this solution into trajectories parameterized by absolute time, the break-points do not need to be chosen uniformly spaced in absolute time.

The constraints represented in (2) can be formally incorporated into an unconstrained optimization problem as follows. We search for a solution to the problem:

$$\min_{\{x_i(s)\}_{i,s}} f^{\text{cost}}(\{x_i(s)\}_{i,s}) + \sum_{s=0}^{\eta-1} \sum_{i>j} f_{r_i, r_j}^{\text{coll}}(x_i(s), x_i(s+1), x_j(s), x_j(s+1)), \quad (3)$$

where $\{x_i(0)\}_p$ and $\{x_i(\eta)\}_p$ are constants rather than optimization variables, and where the function f^{cost} is a function that represents some cost to be minimized (e.g. the integrated kinetic energy or the maximum velocity over all the agents) and the function $f_{r,r'}^{\text{coll}}$ is defined as,

$$f_{r,r'}^{\text{coll}}(\underline{x}, \bar{x}, \underline{x}', \bar{x}') = \mathbb{J}(\|\alpha(\bar{x} - \bar{x}') + (1-\alpha)(\underline{x} - \underline{x}')\| \geq r + r' \forall \alpha \in [0, 1]). \quad (4)$$

In this section, \underline{x} and \bar{x} represent the position of an arbitrary agent of radius r at two consecutive break-points and \underline{x}' and \bar{x}' the position of a second arbitrary agent of radius r' at the same break-points. In the expression above $\mathbb{J}(\cdot)$ takes the value 0 whenever its argument, a clause, is true and takes the value $+\infty$ otherwise. Intuitively, we pay an infinite cost in $f_{r,r'}^{\text{coll}}$ whenever there is a collision, and we pay zero otherwise.

In (3) we can set $f^{\text{cost}}(\cdot)$, to enforce a preference for trajectories satisfying specific properties. For example, we might prefer trajectories for which the total kinetic energy spent by the set of agents is small. In this case, defining $f_C^{\text{cost}}(\underline{x}, \bar{x}) = C\|\underline{x} - \bar{x}\|^2$, we have,

$$f^{\text{cost}}(\{x_i(s)\}_{i,s}) = \frac{1}{p\eta} \sum_{i=1}^p \sum_{s=0}^{\eta-1} f_{C_{i,s}}^{\text{cost}}(x_i(s), x_i(s+1)). \quad (5)$$

where the coefficients $\{C_{i,s}\}$ can account for agents with different masses, different absolute-time intervals between-break points or different preferences regarding which agents we want to be less active and which agents are allowed to move faster.

More simply, we might want to exclude trajectories in which agents move faster than a certain amount, but without distinguishing among all remaining trajectories. For this case we can write,

$$f_C^{\text{cost}}(\underline{x}, \bar{x}) = \mathbb{J}(\|\bar{x} - \underline{x}\| \leq C). \quad (6)$$

In this case, associating each break-point to a time instant, the coefficients $\{C_{i,s}\}$ in expression (5) would represent different limits on the velocity of different agents between different sections of the trajectory. If we want to force all agents to have a minimum velocity we can simply reverse the inequality in (6).

²We replaced the strict inequality in the condition for non-collision by a simple inequality “ \geq ” to avoid technicalities in formulating the optimization problem. Since the agents are round, this allows for a single point of contact between two agents and does not reduce practical relevance.

3.2 Formulation as unconstrained optimization with static obstacles

In many scenarios agents should also avoid collisions with static obstacles. Given two points in space, x_L and x_R , we can forbid all agents from crossing the line segment from x_L to x_R by adding the following term to the function (3): $\sum_{i=1}^p \sum_{s=0}^{\eta-1} f_{x_L, x_R, r_i}^{\text{wall}}(x_i(s), x_i(s+1))$. We recall that r_i is the radius of agent i and

$$f_{x_L, x_R, r}^{\text{wall}}(\underline{x}, \bar{x}) = \mathbb{J}(\|(\alpha \bar{x} + (1 - \alpha)\underline{x}) - (\beta x_R + (1 - \beta)x_L)\| \geq r \text{ for all } \alpha, \beta \in [0, 1]). \quad (7)$$

Notice that f^{coll} can be expressed using f^{wall} . In particular,

$$f_{r, r'}^{\text{coll}}(\underline{x}, \bar{x}, \underline{x}', \bar{x}') = f_{0, 0, r+r'}^{\text{wall}}(\underline{x}' - \underline{x}, \bar{x}' - \bar{x}). \quad (8)$$

We use this fact later to express the minimizer associated with agent-agent collisions using the minimizer associated with agent-obstacle collisions.

When agents move in the plane, i.e. $x_i(s) \in \mathbb{R}^2$ for all $i \in [p]$ and $s+1 \in [\eta+1]$, being able to avoid collisions with a general static line segment allows to automatically avoid collisions with multiple static obstacles of arbitrary polygonal shape. Our numerical experiments only consider agents in the plane and so, in this paper, we only describe the minimizer block for wall collision for a 2D world. In higher dimensions, different obstacle primitives need to be considered.

3.3 Message-passing formulation

To solve (3) using the TWA, we need to specify the topology of the bipartite graph associated with the unconstrained formulation (3) and the operation performed by every minimizer, i.e. the $\overleftarrow{\rho}$ -weight update logic and x -variable update equations. We postpone describing the choice of initial values and internal parameters until Section 4.

We first describe the bipartite graph. To be concrete, let us assume that the cost functional has the form of (5). The unconstrained formulation (3) then tells us that the global objective function is the sum of $\eta p(p+1)/2$ terms: $\eta p(p-1)/2$ functions f^{coll} and ηp functions f_C^{cost} . These functions involve a total of $(\eta+1)p$ variables out of which only $(\eta-1)p$ are free (since the initial and final configurations are fixed). Correspondingly, the bipartite graph along which messages are passed has $\eta p(p+1)/2$ minimizer-nodes that connect to the $(\eta+1)p$ equality-nodes. In particular, the equality-node associated with the break-point variable $x_i(s)$, $\eta > s > 0$, is connected to $2(p-1)$ different g^{coll} minimizer-nodes and two different g_C^{cost} minimizer-nodes. If $s=0$ or $s=\eta$ the equality-node only connects to half as many g^{coll} nodes and g_C^{cost} nodes.

We now describe the different minimizers. Every minimizer basically is a special case of (1).

3.3.1 Agent-agent collision minimizer

We start with the minimizer associated with the functions f^{coll} , that we denoted by g^{coll} . This minimizer receives as parameters the radius, r and r' , of the two agents whose collision it is avoiding. The minimizer takes as input a set of incoming n -messages, $\{\underline{n}, \bar{n}, \underline{n}', \bar{n}'\}$, and associated $\overleftarrow{\rho}$ -weights, $\{\overleftarrow{\rho}, \overleftarrow{\rho}, \overleftarrow{\rho}', \overleftarrow{\rho}'\}$, and outputs a set of updated x -variables according to expression (9). Messages \underline{n} and \bar{n} come from the two equality-nodes associated with the positions of one of the agents at two consecutive break-points and \underline{n}' and \bar{n}' from the corresponding equality-nodes for the other agent.

$$g^{\text{coll}}(\underline{n}, \bar{n}, \underline{n}', \bar{n}', \overleftarrow{\rho}, \overleftarrow{\rho}, \overleftarrow{\rho}', \overleftarrow{\rho}', r, r') = \arg \min_{\{\underline{x}, \bar{x}, \underline{x}', \bar{x}'\}} f_{r, r'}^{\text{coll}}(\underline{x}, \bar{x}, \underline{x}', \bar{x}') + \frac{\overleftarrow{\rho}}{2} \|\underline{x} - \underline{n}\|^2 + \frac{\overleftarrow{\rho}}{2} \|\bar{x} - \bar{n}\|^2 + \frac{\overleftarrow{\rho}'}{2} \|\underline{x}' - \underline{n}'\|^2 + \frac{\overleftarrow{\rho}'}{2} \|\bar{x}' - \bar{n}'\|^2. \quad (9)$$

The update logic for the weights $\overleftarrow{\rho}$ for this minimizer is simple. If the trajectory from \underline{n} to \bar{n} for an agent of radius r does not collide with the trajectory from \underline{n}' to \bar{n}' for an agent of radius r' then set all the outgoing weights $\overleftarrow{\rho}$ to zero. Otherwise set them all to ρ_0 . The outgoing zero weights indicate to the receiving equality-nodes in the bipartite graph that the collision constraint for this pair of agents is inactive and that the values it receives from this minimizer-node should be ignored when computing the consensus values of the receiving equality-nodes.

The solution to (9) is found using the agent-obstacle collision minimizer that we describe next.

3.3.2 Agent-obstacle collision minimizer

The minimizer for f^{wall} is denoted by g^{wall} . It is parameterized by the obstacle position $\{x_L, x_R\}$ as well as the radius of the agent that needs to avoid the obstacle. It receives two n -messages, $\{\underline{n}, \bar{n}\}$, and corresponding weights $\{\overleftarrow{\rho}, \overleftarrow{\bar{\rho}}\}$, from the equality-nodes associated with two consecutive positions of an agent that needs to avoid the obstacle. Its output, the x -variables, are defined as

$$g^{\text{wall}}(\underline{n}, \bar{n}, r, x_L, x_R, \overleftarrow{\rho}, \overleftarrow{\bar{\rho}}) = \arg \min_{\{\underline{x}, \bar{x}\}} f_{x_L, x_R, r}^{\text{wall}}(\underline{x}, \bar{x}) + \frac{\overleftarrow{\rho}}{2} \|\underline{x} - \underline{n}\|^2 + \frac{\overleftarrow{\bar{\rho}}}{2} \|\bar{x} - \bar{n}\|^2. \quad (10)$$

When agents move in the plane (2D), this minimizer can be solved by reformulating the optimization in (10) as a mechanical problem involving a system of springs that we can solve exactly and efficiently. This reduction is explained in the supplementary material in Section D and the solution to the mechanical problem is explained in Section I.

The update logic for the $\overleftarrow{\rho}$ -weights is similar to that of the g^{coll} minimizer. If an agent of radius r going from \underline{n} and \bar{n} does not collide with the line segment from x_L to x_R then set all outgoing weights to zero because the constraint is inactive; otherwise set all the outgoing weights to ρ_0 .

Notice that, from (8), it follows that the agent-agent minimizer g^{coll} can be expressed using g^{wall} . More concretely, as proved in the supplementary material, Section C,

$$g^{\text{coll}}(\underline{n}, \bar{n}, \underline{n}', \bar{n}', \overleftarrow{\rho}, \overleftarrow{\bar{\rho}}, \overleftarrow{\rho}', \overleftarrow{\bar{\rho}}', r, r') = M_2 g^{\text{wall}} \left(M_1 \cdot \{\underline{n}, \bar{n}, \underline{n}', \bar{n}', \overleftarrow{\rho}, \overleftarrow{\bar{\rho}}, \overleftarrow{\rho}', \overleftarrow{\bar{\rho}}', r, r'\} \right),$$

for a constant rectangular matrix M_1 and a matrix M_2 that depend on $\{\underline{n}, \bar{n}, \underline{n}', \bar{n}', \overleftarrow{\rho}, \overleftarrow{\bar{\rho}}, \overleftarrow{\rho}', \overleftarrow{\bar{\rho}}'\}$.

3.3.3 Minimum energy and maximum (minimum) velocity minimizer

When f^{cost} can be decomposed as in (5), the minimizer associated with the functions f^{cost} is denoted by g^{cost} and receives as input two n -messages, $\{\underline{n}, \bar{n}\}$, and corresponding weights, $\{\overleftarrow{\rho}, \overleftarrow{\bar{\rho}}\}$. The messages come from two equality-nodes associated with two consecutive positions of an agent. The minimizer is also parameterized by a cost factor c . It outputs a set of updated x -messages defined as

$$g^{\text{cost}}(\underline{n}, \bar{n}, \overleftarrow{\rho}, \overleftarrow{\bar{\rho}}, c) = \arg \min_{\{\underline{x}, \bar{x}\}} f_c^{\text{cost}}(\underline{x}, \bar{x}) + \frac{\overleftarrow{\rho}}{2} \|\underline{x} - \underline{n}\|^2 + \frac{\overleftarrow{\bar{\rho}}}{2} \|\bar{x} - \bar{n}\|^2. \quad (11)$$

The update logic for the $\overleftarrow{\rho}$ -weights of the minimum energy minimizer is very simple: always set all outgoing weights $\overleftarrow{\rho}$ to ρ_0 . The update logic for the $\overleftarrow{\bar{\rho}}$ -weights of the maximum velocity minimizer is the following. If $\|\underline{n} - \bar{n}\| \leq c$ set all outgoing weights to zero. Otherwise, set them to ρ_0 . The update logic for the minimum velocity minimizer is similar. If $\|\underline{n} - \bar{n}\| \geq c$, set all the $\overleftarrow{\bar{\rho}}$ -weights to zero. Otherwise set them to ρ_0 .

The solution to the minimum energy, maximum velocity and minimum velocity minimizer is written in the supplementary material in Sections E, F, and G respectively.

4 Numerical results

We now report on the performance of our algorithm (see Appendix J for an important comment on the anytime properties of our algorithm). Note that the lack of open-source scalable algorithms for global trajectory planning in the literature makes it difficult to benchmark our performance against other methods. Also, in a paper it is difficult to appreciate the gracefulness of the discovered trajectory optimizations, so we include a video in the supplementary material that shows final optimized trajectories as well as intermediate results as the algorithm progresses for a variety of additional scenarios, including those with obstacles. All the tests described here are for agents in a two-dimensional plane. All tests but the last were performed using six cores of a 3.4GHz i7 CPU.

The different tests did not require any special tuning of parameters. In particular, the step-size in [13] (their α variable) is always 0.1. In order to quickly equilibrate the system to a reasonable set of variables and to wash out the importance of initial conditions, the default weight ρ_0 was set equal to a small value ($\eta p \times 10^{-5}$) for the first 20 iterations and then set to 1 for all further iterations.

The first test considers scenario CONF1: p (even) agents of radius r , equally spaced around on a circle of radius R , are each required to exchange position with the corresponding antipodal agent, $r = (5/4)R \sin(\pi/2(p - 4))$. This is a classical difficult test scenario because the straight line motion of all agents to their goal would result in them all colliding in the center of the circle. We compare the convergence time of the TWA with a similar version using standard ADMM to perform the optimizations. In this test, the algorithm’s initial value for each variable in the problem was set to the corresponding initial position of each agent. The objective is to minimize the total kinetic energy (C in the energy minimizer is set to 1). Figure 2-left shows that the TWA scales better with p than classic ADMM and typically gives an order of magnitude speed-up. Please see Appendix K for a further comment on the scaling of the convergence time of ADMM and TWA with p .




Figure 2: Left: Convergence time using standard ADMM (dashed lines) and using TWA (solid lines). Middle: Distribution of total energy and time for convergence with random initial conditions ($p = 20$ and $\eta = 5$). Right: Convergence time using a different number of cores ($\eta = 5$).

The second test for CONF1 analyzes the sensitivity of the convergence time and objective value when the variables’ value at the first iteration are chosen uniformly at random in the smallest space-time box that includes the initial and final configuration of the robots. Figure 2-middle shows that, although there is some spread on the convergence time, our algorithm seems to reliably converge to relatively similar-cost local minima (other experiments show that the objective value of these minima is around 5 times smaller than that found when the algorithm is run using only the collision avoidance minimizers without a kinetic energy cost term). As would be expected, the precise trajectories found vary widely between different random runs.

Still for CONF1, and fixed initial conditions, we parallelize our method using several cores of a 2.66GHz i7 processor and a very primitive scheduling/synchronization scheme. Although this scheme does not fully exploit parallelization, Figure 2-right does show a speed-up as the number of cores increases and the larger p is, the greater the speed-up. We stall when we reach the twelve physical cores available and start using virtual cores.

Finally, Figure 3-left compares the convergence time to optimize the total energy with the time to simply find a feasible (i.e. collision-free) solution. The agents initial and final configuration is randomly chosen in the plane (CONF2). Error bars indicate \pm one standard deviation. Minimizing the kinetic energy is orders of magnitude computationally more expensive than finding a feasible solution, as is clear from the different magnitude of the left and right scale of Figure 3-left.




Figure 3: Left: Convergence time when minimizing energy (blue scale/dashed lines) and to simply find a feasible solution (red scale/solid lines). Right: (For Section 5). Convergence-time distribution for each epoch using our method (blue bars) and using the MIQP of [12] (red bars and star-values).

5 Local trajectory planning based on velocity obstacles

In this section we show how the joint optimization presented in [12], which is based on the concept of velocity obstacles [11] (VO), can be also solved via the message-passing TWA. In VO, given the current position $\{x_i(0)\}_{i \in [p]}$ and radius $\{r_i\}$ of all agents, a new velocity command is computed jointly for all agents minimizing the distance to their preferred velocity $\{v_i^{\text{ref}}\}_{i \in [p]}$. This new velocity command must guarantee that the trajectories of all agents remain collision-free for at least a time horizon τ . New collision-free velocities are computed every $\alpha\tau$ seconds, $\alpha < 1$, until all agents reach their final configuration. Following [12], and assuming an obstacle-free environment and first order dynamics, the collision-free velocities are given by,

$$\underset{\{v_i\}_{i \in [p]}}{\text{minimize}} \sum_{i \in [p]} C_i \|v_i - v_i^{\text{ref}}\|^2 \text{ s.t. } \|(x_i(0) + v_i t) - (x_j(0) + v_j t)\| \geq r_i + r_j \quad \forall i \in [p], t \in [0, \tau].$$

Since the velocities $\{v_i\}_{i \in [p]}$ are related linearly to the final position of each object after τ seconds, $\{x_i(\tau)\}_{i \in [p]}$, a simple change of variables allows us to reformulate the above problem as,

$$\underset{\{x_i\}_{i \in [p]}}{\text{minimize}} \sum_{i \in [p]} C'_i \|x_i - x_i^{\text{ref}}\|^2$$

$$\text{ s.t. } \|(1 - \alpha)(x_i(0) - x_j(0)) + \alpha(x_i - x_j)\| \geq r_i + r_j \quad \forall j > i \in [p], \alpha \in [0, 1] \quad (12)$$

where $C'_i = C_i/\tau^2$, $x_i^{\text{ref}} = x_i(0) + v_i^{\text{ref}}\tau$ and we have dropped the τ in $x_i(\tau)$. The above problem, extended to account for collisions with the static line segments $\{x_{Rk}, x_{Lk}\}_k$, can be formulated in an unconstrained form using the functions f^{cost} , f^{coll} and f^{wall} . Namely,

$$\min_{\{x_i\}_i} \sum_{i \in [p]} f_{C'_i}^{\text{cost}}(x_i, x_i^{\text{ref}}) + \sum_{i > j} f_{r_i, r_j}^{\text{coll}}(x_i(0), x_i, x_j(0), x_j) + \sum_{i \in [p]} \sum_k f_{x_{Rk}, x_{Lk}, r_i}^{\text{wall}}(x_i(0), x_i). \quad (13)$$

Note that $\{x_i(0)\}_i$ and $\{x_i^{\text{ref}}\}_i$ are constants, not variables being optimized. Given this formulation, the TWA can be used to solve the optimization. All corresponding minimizers are special cases of minimizers derived in the previous section for global trajectory planning (see Section H in the supplementary material for details). Figure 3-right shows the distribution of the time to solve (12) for CONF1. We compare the mixed integer quadratic programming (MIQP) approach from [12] with ours. Our method finds a local minima of exactly (13), while [12] finds a global minima of an approximation to (13). Specifically, [12] requires approximating the search domain by hyperplanes and an additional branch-and-bound algorithm while ours does not. Both approaches use a mechanism for breaking the symmetry from CONF1 and avoid deadlocks: theirs uses a preferential rotation direction for agents, while we use agents with slightly different C coefficients in their energy minimizers ($C_{i^{\text{th}} \text{ agent}} = 1 + 0.001i$). Both simulations were done on a single 2.66GHz core. The results show the order of magnitude is similar, but, because our implementation is done in Java while [12] uses Matlab-mex interface of CPLEX 11, the results are not exactly comparable.

6 Conclusion and future work

We have presented a novel algorithm for global and local planning of the trajectory of multiple distinct agents, a problem known to be hard. The solution is based on solving a non-convex optimization problem using TWA, a modified ADMM. Its similarity to ADMM brings scalability and easy parallelization. However, using TWA improves performance considerably. Our implementation of the algorithm in Java on a regular desktop computer, using a basic scheduler/synchronization over its few cores, already scales to hundreds of agents and achieves real-time performance for local planning.

The algorithm can flexibly account for obstacles and different cost functionals. For agents in the plane, we derived explicit expressions that account for static obstacles, moving obstacles, and dynamic constraints on the velocity and energy. Future work should consider other restrictions on the smoothness of the trajectory (e.g. acceleration constraints) and provide fast solvers to our minimizers for agents in 3D.

The message-passing nature of our algorithm hints that it might be possible to adapt our algorithm to do planning in a decentralized fashion. For example, minimizers like g^{coll} could be solved by message exchange between pairs of agents within a maximum communication radius. It is an open problem to build a practical communication-synchronization scheme for such an approach.

References

- [1] Javier Alonso-Mora, Andreas Breitenmoser, Martin Rufli, Roland Siegwart, and Paul Beardsley. Image and animation display with multiple mobile robots. *31(6):753–773*, 2012.
- [2] Peter R. Wurman, Raffaello D’Andrea, and Mick Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–19, 2008.
- [3] Stephen J. Guy, Jatin Chhugani, Changkyu Kim, Nadathur Satish, Ming Lin, Dinesh Manocha, and Pradeep Dubey. Clearpath: highly parallel collision avoidance for multi-agent simulation. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 177–187, 2009.
- [4] John H. Reif. Complexity of the mover’s problem and generalizations. In *IEEE Annual Symposium on Foundations of Computer Science*, pages 421–427, 1979.
- [5] John E. Hopcroft, Jacob T. Schwartz, and Micha Sharir. On the complexity of motion planning for multiple independent objects; pspace-hardness of the ”warehouseman’s problem”. *The International Journal of Robotics Research*, 3(4):76–88, 1984.
- [6] Maren Bennewitz, Wolfram Burgard, and Sebastian Thrun. Finding and optimizing solvable priority schemes for decoupled path planning techniques for teams of mobile robots. *Robotics and Autonomous Systems*, 41(2–3):89–99, 2002.
- [7] Daniel Mellinger, Alex Kushleyev, and Vijay Kumar. Mixed-integer quadratic program trajectory generation for heterogeneous quadrotor teams. In *IEEE International Conference on Robotics and Automation*, pages 477–483, 2012.
- [8] Federico Augugliaro, Angela P. Schoellig, and Raffaello D’Andrea. Generation of collision-free trajectories for a quadcopter fleet: A sequential convex programming approach. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1917–1922, 2012.
- [9] Steven M. LaValle and James J. Kuffner. Randomized kinodynamic planning. *The International Journal of Robotics Research*, 20(5):378–400, 2001.
- [10] Oussama Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98, 1986.
- [11] Paolo Fiorini and Zvi Shiller. Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research*, 17(7):760–772, 1998.
- [12] Javier Alonso-Mora, Martin Rufli, Roland Siegwart, and Paul Beardsley. Collision avoidance for multiple agents with joint utility maximization. In *IEEE International Conference on Robotics and Automation*, 2013.
- [13] Nate Derbinsky, José Bento, Veit Elser, and Jonathan S. Yedidia. An improved three-weight message-passing algorithm. *arXiv:1305.1961 [cs.AI]*, 2013.
- [14] G. David Forney Jr. Codes on graphs: Normal realizations. *IEEE Transactions on Information Theory*, 47(2):520–548, 2001.
- [15] Sertac Karaman and Emilio Frazzoli. Incremental sampling-based algorithms for optimal motion planning. *arXiv preprint arXiv:1005.0416*, 2010.
- [16] R. Glowinski and A. Marrocco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité, d’une class de problèmes de Dirichlet non linéaire. *Revue Française d’Automatique, Informatique, et Recherche Opérationnelle*, 9(2):41–76, 1975.
- [17] Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976.
- [18] Hugh Everett III. Generalized lagrange multiplier method for solving problems of optimum allocation of resources. *Operations Research*, 11(3):399–417, 1963.
- [19] Magnus R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4(5):303–320, 1969.
- [20] Magnus R. Hestenes. Multiplier and gradient methods. In L.A. Zadeh et al., editor, *Computing Methods in Optimization Problems 2*. Academic Press, New York, 1969.
- [21] M.J.D. Powell. A method for nonlinear constraints in minimization problems. In R. Fletcher, editor, *Optimization*. Academic Press, London, 1969.
- [22] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

Supplementary material for “A message-passing algorithm for multi-agent trajectory planning”

This document gives details on the message-passing algorithm we use and how to implement all minimizers described in the paper. It also includes some more comments on our numerical results and related literature. A short movie clip showing the behaviour of our algorithm can be found at <http://youtu.be/youGckVT8Bew>

A Comment on related literature

A*-search based methods and sampling-based methods require exploring a continuous domain using discrete graph structures. For problems with many degrees of freedom or complex kinematic and dynamic constraints, as when dealing with multiple agents or manipulators, fixed-grid search methods are impractical. Alternatively, exploration can be done using sampling algorithms with proved asymptotic convergence to the optimal solution [15]. However, as the dimensionality of the configuration space increases, the convergence rate degrades and the local planners required by the exploration loop become harder to implement. In addition, as pointed out in [9], sampling algorithms cannot easily produce solutions where multiple agents move in tight spaces, like in CONF1 with obstacles. Some of the disadvantages of using discrete random search structures are even visible in extremely simple scenarios. For example, for a single holonomic agent that needs to move as quickly as possible between two points in free-space, [15] require around 10000 samples on their RRT* method to find something close to the shortest-path solution. For our algorithm this is a trivial scenario: it outputs the optimal straight-line solution in 200 iterations and 37 msec. in our Java implementation.

B Full description of the improved three-weight message-passing algorithm of [13]

First we give a self-contained (complete) description of the three-weight algorithm TWA from [13]. Their method is an improvement of the alternating direction method of multipliers (ADMM)³. Assume we want to solve

$$\min_{x \in \mathbb{R}^d} \sum_{b=1}^l f_b(x_{\partial b}), \quad (14)$$

where the set $x_{\partial b} = \{x_j : j \in \partial b\}$ is a vector obtained by considering the subset of entries of x with index in $\partial b \subseteq [d]$. The functions f_b do not need to be convex or smooth for the algorithm to be well-defined. However, the algorithm is only guaranteed to find the global minimum under convexity [22].

Start by forming the following bipartite graph consisting of minimizer-nodes and equality-nodes. Create one minimizer-node, labeled “ g ”, per function f_b and one equality-node, labeled “ $=$ ”, per variable x_j . There are l minimizer-nodes and d equality-nodes in total. If function f_b depends on variable x_j , create an edge (b, j) connecting b and j (see Figure B for a general representation).

The algorithm in [13] works by repetitively updating seven kind of variables. These can be listed as follows. Every equality-node j has a corresponding variable z_j . Every edge (b, j) from minimizer-node b to equality-node j has a corresponding variable $x_{b,j}$, variable $u_{b,j}$, message $n_{b,j}$, message $m_{b,j}$, weight $\vec{\rho}_{b,j}$ and weight $\overleftarrow{\rho}_{b,j}$.

To start the method, one specifies the initial values $\{z_j^0\}$, $\{u_{b,j}^0\}$ and $\{\overleftarrow{\rho}_{b,j}^0\}$. Then, at every iteration k , repeat the following.

³ADMM is a decomposition procedure for solving optimization problems. It coordinates the solutions to small local sub-problems to solve a large global problem. Hence, it is useful to derive parallel algorithms. It was introduced in [16] and [17] but is closely related to previous work as the dual decomposition method introduced by [18] and the method of multipliers introduced by [19, 20] and [21]. For a good review on ADMM see [22], where you can also find a self-contained proof of its convergence for convex problems.




Figure 4: Bipartite graph used by the message-passing algorithm to solve (14). The algorithm works by updating seven kind of variables, also shown in picture.

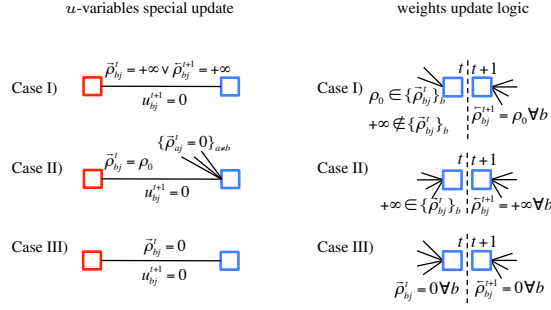


Figure 5: Left: Special update-rules for variables u . Right: Update rule for variables $\bar{\rho}$.

1. Construct the n -message for every edge (b, j) as $n_{b,j}^k = z_j^k - u_{b,j}^k$.
2. Update the x -variables for every edge (b, j) . All x -variables associated to edges incident on the same minimizer-node b , i.e. $\{x_{b,j}^k\}_{j \in \partial b}$, are updated simultaneously by the minimizer g_b associated to the function f_b ,

$$\begin{aligned} \{x_{b,j}^{k+1}\}_j &= g_b(\{n_{b,j}^k\}_j, \{\bar{\rho}_{b,j}^k\}_j) \\ &\equiv \arg \min_{\{\tilde{x}_{b,j}^k\}_j} f_b(\{\tilde{x}_{b,j}^k\}_j) + \frac{1}{2} \sum_j \bar{\rho}_{b,j}^k (\tilde{x}_{b,j}^k - n_{b,j}^k)^2. \end{aligned}$$

In the expression above we write $\{\}_j$ and \sum_j for $\{j \in \partial b$ and $\sum_{j \in \partial b}$ respectively. If the $\arg \min$ returns a set with more than one element, choose a value uniformly at random from this set.

3. Compute the outgoing weights for all edges (b, j) . All weights associated to edges leaving the same minimizer-node b , i.e. $\{\bar{\rho}_{b,j}^k\}_{j \in \partial b}$, are updated simultaneously according to a user-defined logic that should depend on the problem defined by equation (14). This logic can assign three possible values for the weights: $\bar{\rho}_{b,j}^k = 0$, $\bar{\rho}_{b,j}^k = \rho_0$ or $\bar{\rho}_{b,j}^k = \infty$, where $\rho_0 > 0$ is some pre-specified constant. The idea is to use these three values to inform the equality-nodes of how certain the minimizer-node b is that the current variable $x_{b,j}^k$ is the optimal value for variable x_j in equation (14). A weight of ∞ is used to signal total certainty, 0 for no certainty at all and ρ_0 for all the other scenarios. Later, equality-nodes average the information coming from the minimizer-nodes by these weights to update the consensus variables z .

4. Construct the m -messages for every edge (b, j) as $m_{b,j}^k = x_{b,j}^{k+1} + u_{b,j}^k$.

5. Update the z -variable for every equality-node j as

$$z_j^{k+1} = \frac{\sum_{b \in \partial j} \bar{\rho}_{b,j}^k m_{b,j}^k}{\sum_{b \in \partial j} \bar{\rho}_{b,j}^k}.$$

The set ∂j contains all the minimizer-nodes that connect to equality-node j . If all weights $\{\overrightarrow{\rho}_{b,j}^k\}_{b \in \partial j}$ are zero in the previous expression, treat them as 1.

6. Compute the updated weights $\overleftarrow{\rho}^{k+1}$. Weights leaving the same equality-node j , i.e. $\{\overleftarrow{\rho}_{b,j}^{k+1}\}_{b \in \partial j}$, are computed simultaneously. The update logic is described in Figure 5-right. According to which of the three distinct scenarios $\overrightarrow{\rho}^k$ falls, $\overleftarrow{\rho}^{k+1}$ is uniquely determined. This logic again assigns three possible values for the weights $\{0, \rho_0, \infty\}$.
7. Update the u -variables for all edges. If edge (b, j) has $\overrightarrow{\rho}_{b,j}^k = \rho_0$ and $\overleftarrow{\rho}_{b,j}^{k+1} = \rho_0$ then $u_{b,j}^{k+1} = u_{b,j}^k + (\alpha/\rho_0)(x_{b,j}^{k+1} - z_j^{k+1})$, where α is a pre-specified constant. Otherwise, choose $u_{b,j}^{k+1}$ according to the three scenarios described in Figure 5-left, depending on the weights $\overrightarrow{\rho}$ and $\overleftarrow{\rho}$.

In short, given ρ_0 , α , the initial values $\{z_j^0\}$, $\{u_{b,j}^0\}$ and $\{\overleftarrow{\rho}_{b,j}^0\}$, all the minimizers $\{g_b\}$, with corresponding update logic for $\overrightarrow{\rho}$, and the bipartite graph, the method is completely specified. If all weights, $\overrightarrow{\rho}$ and $\overleftarrow{\rho}$, are set to ρ_0 across all iterations, the described method reduces to classical ADMM, interpreted as a message-passing algorithm. Finally, notice that at each time step k , all the variables associated with each edge can be updated in parallel. In particular, the update of the x -variables, usually the most expensive operation, can be parallelized.

C Agent-agent collision minimizer

Here we give the details of how to write the agent-agent collision minimizer, g^{coll} using the agent-obstacle minimizer g^{wall} of equation (10).

First recall that $f_{r,r'}^{\text{coll}}(\underline{x}, \overline{x}, \underline{x}', \overline{x}') = f_{0,0,r+r'}^{\text{wall}}(\underline{x}' - \underline{x}, \overline{x}' - \overline{x})$. Then, rewrite (9) as,

$$g^{\text{coll}}(\underline{n}, \overline{n}, \underline{n}', \overline{n}', \overleftarrow{\rho}, \overleftarrow{\rho}', \overleftarrow{\rho}', r, r') = \arg \min_{\{\underline{x}, \overline{x}, \underline{x}', \overline{x}'\}} \left[f_{0,0,r+r'}^{\text{wall}}(\underline{x}' - \underline{x}, \overline{x}' - \overline{x}) + \frac{\overleftarrow{\rho}}{2} \|\underline{x} - \underline{n}\|^2 + \frac{\overleftarrow{\rho}}{2} \|\overline{x} - \overline{n}\|^2 + \frac{\overleftarrow{\rho}'}{2} \|\underline{x}' - \underline{n}'\|^2 + \frac{\overleftarrow{\rho}'}{2} \|\overline{x}' - \overline{n}'\|^2 \right]. \quad (15)$$

Now introduce the following variables $\underline{v} = \underline{x}' - \underline{x}$, $\underline{u} = \underline{x}' + \underline{x}$, $\overline{v} = \overline{x}' - \overline{x}$ and $\overline{u} = \overline{x}' + \overline{x}$. The function being minimized in (15) can be written as

$$f_{0,0,r+r'}^{\text{wall}}(\underline{v}, \overline{v}) + \frac{\overleftarrow{\rho}}{2} \left\| \frac{\underline{u} - \underline{v}}{2} - \underline{n} \right\|^2 + \frac{\overleftarrow{\rho}}{2} \left\| \frac{\overline{u} - \overline{v}}{2} - \overline{n} \right\|^2 + \frac{\overleftarrow{\rho}'}{2} \left\| \frac{\underline{u} + \underline{v}}{2} - \underline{n}' \right\|^2 + \frac{\overleftarrow{\rho}'}{2} \left\| \frac{\overline{u} + \overline{v}}{2} - \overline{n}' \right\|^2. \quad (16)$$

Now notice that we can write,

$$\begin{aligned} & \frac{\overleftarrow{\rho}}{2} \left\| \frac{\underline{u} - \underline{v}}{2} - \underline{n} \right\|^2 + \frac{\overleftarrow{\rho}'}{2} \left\| \frac{\underline{u} + \underline{v}}{2} - \underline{n}' \right\|^2 = \frac{\overleftarrow{\rho}}{8} \|\underline{u} - \underline{v} - 2\underline{n}\|^2 + \frac{\overleftarrow{\rho}'}{8} \|\underline{u} + \underline{v} - 2\underline{n}'\|^2 \\ & = \frac{\overleftarrow{\rho}}{8} (\|\underline{u}\|^2 + \|\underline{v} + 2\underline{n}\|^2 - 2\langle \underline{u}, \underline{v} + 2\underline{n} \rangle) + \frac{\overleftarrow{\rho}'}{8} (\|\underline{u}\|^2 + \|\underline{v} - 2\underline{n}'\|^2 + 2\langle \underline{u}, \underline{v} - 2\underline{n}' \rangle) \\ & = \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{8} \|\underline{u}\|^2 + 2 \left\langle \underline{u}, \frac{\overleftarrow{\rho}' - \overleftarrow{\rho}}{8} \underline{v} - \frac{\overleftarrow{\rho} \underline{n} + \overleftarrow{\rho}' \underline{n}'}{4} \right\rangle + \frac{\overleftarrow{\rho}}{8} \|\underline{v} + 2\underline{n}\|^2 + \frac{\overleftarrow{\rho}'}{8} \|\underline{v} - 2\underline{n}'\|^2 \\ & = \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{8} \left\| \underline{u} - \left(\frac{\overleftarrow{\rho} - \overleftarrow{\rho}'}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \underline{v} + \frac{2(\overleftarrow{\rho} \underline{n} + \overleftarrow{\rho}' \underline{n}')}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \right) \right\|^2 + \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{8} \left\| \underline{v} - \frac{2(\overleftarrow{\rho}' \underline{n}' - \overleftarrow{\rho} \underline{n})}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \right\|^2 \\ & + C(\underline{n}, \overleftarrow{\rho}, \underline{n}', \overleftarrow{\rho}'), \end{aligned}$$

where $C(\underline{n}, \overleftarrow{\rho}, \underline{n}', \overleftarrow{\rho}')$ is a constant that depends on the variables $\{\underline{n}, \overleftarrow{\rho}, \underline{n}', \overleftarrow{\rho}'\}$. A similar manipulation can be done to $\frac{\overleftarrow{\rho}}{2} \left\| \frac{\overline{u} - \overline{v}}{2} - \overline{n} \right\|^2 + \frac{\overleftarrow{\rho}'}{2} \left\| \frac{\overline{u} + \overline{v}}{2} - \overline{n}' \right\|^2$. Therefore, the expression (16) can be

rewritten as

$$\begin{aligned}
f_{0,0,r+r'}^{\text{wall}}(\underline{v}, \bar{v}) &+ \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{8} \left\| \underline{v} - \frac{2(\overleftarrow{\rho}' \underline{n}' - \overleftarrow{\rho} \underline{n})}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \right\|^2 + \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{8} \left\| \bar{v} - \frac{2(\overleftarrow{\rho}' \bar{n}' - \overleftarrow{\rho} \bar{n})}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \right\|^2 \\
&+ \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{8} \left\| \underline{u} - \left(\frac{\overleftarrow{\rho} - \overleftarrow{\rho}'}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \underline{v} + \frac{2(\overleftarrow{\rho} \underline{n} + \overleftarrow{\rho}' \underline{n}')}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \right) \right\|^2 \\
&+ \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{8} \left\| \bar{u} - \left(\frac{\overleftarrow{\rho} - \overleftarrow{\rho}'}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \bar{v} + \frac{2(\overleftarrow{\rho} \bar{n} + \overleftarrow{\rho}' \bar{n}')}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \right) \right\|^2 + C(\underline{n}, \underline{n}', \bar{n}, \bar{n}', \overleftarrow{\rho}, \overleftarrow{\rho}', \overleftarrow{\rho}, \overleftarrow{\rho}'), \quad (17)
\end{aligned}$$

where $C(\underline{n}, \underline{n}', \bar{n}, \bar{n}', \overleftarrow{\rho}, \overleftarrow{\rho}', \overleftarrow{\rho}, \overleftarrow{\rho}')$ is a constant that depends on the variables $\underline{n}, \underline{n}', \bar{n}, \bar{n}', \overleftarrow{\rho}, \overleftarrow{\rho}'$, $\overleftarrow{\rho}$ and $\overleftarrow{\rho}'$.

Let $\{\underline{v}^*, \bar{v}^*, \underline{u}^*, \bar{u}^*\}$ be a set of values that minimizes equation (17). We have,

$$\{\underline{v}^*, \bar{v}^*\} \in g^{\text{wall}} \left(\frac{2(\overleftarrow{\rho}' \underline{n}' - \overleftarrow{\rho} \underline{n})}{\overleftarrow{\rho} + \overleftarrow{\rho}'}, \frac{2(\overleftarrow{\rho}' \bar{n}' - \overleftarrow{\rho} \bar{n})}{\overleftarrow{\rho} + \overleftarrow{\rho}'}, r + r', 0, 0, \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{4}, \frac{\overleftarrow{\rho} + \overleftarrow{\rho}'}{4} \right), \quad (18)$$

$$\{\underline{u}^*, \bar{u}^*\} = \left\{ \frac{\overleftarrow{\rho} - \overleftarrow{\rho}'}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \underline{v}^* + \frac{2(\overleftarrow{\rho} \underline{n} + \overleftarrow{\rho}' \underline{n}')}{\overleftarrow{\rho} + \overleftarrow{\rho}'}, \frac{\overleftarrow{\rho} - \overleftarrow{\rho}'}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \bar{v}^* + \frac{2(\overleftarrow{\rho} \bar{n} + \overleftarrow{\rho}' \bar{n}')}{\overleftarrow{\rho} + \overleftarrow{\rho}'} \right\}. \quad (19)$$

We can now produce a set of values that satisfy

$$\{\underline{x}^*, \bar{x}^*, \underline{x}'^*, \bar{x}'^*\} \in g^{\text{coll}}(\underline{n}, \bar{n}, \underline{n}', \bar{n}', \overleftarrow{\rho}, \overleftarrow{\rho}, \overleftarrow{\rho}', \overleftarrow{\rho}', r, r')$$

using the following relation,

$$\{\underline{x}^*, \bar{x}^*, \underline{x}'^*, \bar{x}'^*\} = \left\{ \frac{\underline{u}^* - \underline{v}^*}{2}, \frac{\bar{u}^* - \bar{v}^*}{2}, \frac{\underline{v}^* + \underline{u}^*}{2}, \frac{\bar{v}^* + \bar{u}^*}{2} \right\}.$$

In fact, all values $\{\underline{x}^*, \bar{x}^*, \underline{x}'^*, \bar{x}'^*\} \in g^{\text{coll}}(\underline{n}, \bar{n}, \underline{n}', \bar{n}', \overleftarrow{\rho}, \overleftarrow{\rho}, \overleftarrow{\rho}', \overleftarrow{\rho}', r, r')$ can be obtained from some $\{\underline{v}^*, \bar{v}^*, \underline{u}^*, \bar{u}^*\}$ that minimizes equation (17). In other words, the minimizer g^{coll} can be expressed in terms of the minimizer g^{wall} by means of a linear transformation.

Minimizers can receive zero-weight messages $\overleftarrow{\rho}$ from their neighboring equality nodes. In (18) and (19), this can lead to indeterminacies. We address this as follows. If $\overleftarrow{\rho}$ and $\overleftarrow{\rho}'$ are simultaneously zero then we compute (18) and (19) in the limit when $\overleftarrow{\rho} = \overleftarrow{\rho}' \rightarrow 0^+$. When implementing this on software, we simply replace them by small equal values. The fact that $\overleftarrow{\rho} = \overleftarrow{\rho}'$ resolves the indeterminacies in the fractions and taking the limit to zero from above guarantees that the wall minimizer g^{wall} , that is solved using a mechanical analogy involving springs, is well behaved (See Section I). If $\overleftarrow{\rho}$ and $\overleftarrow{\rho}'$ are simultaneously zero, we perform a similar operation.

D Agent-obstacle collision minimizer

In Section C we expressed the agent-agent collision minimizer by applying a linear transformation to the agent-obstacle collision minimizer. Now we show how the agent-obstacle minimizer can be posed as a classical mechanical problem involving a system of springs. Although the relationship in Section C holds in general, the transformation presented in this section holds only when the agents move in the plane, i.e. $x_i(s) \in \mathbb{R}^2 \forall s, i$. Similar transformations should hold in higher dimensions.

When the obstacle is a line-segment $[x_L, x_R]$, the agent-obstacle minimizer (10) solves the following non-convex optimization problem,

$$\underset{\{\underline{x}, \bar{x}\}}{\text{minimize}} \left[\frac{\overleftarrow{\rho}}{2} \|\underline{x} - \underline{n}\|^2 + \frac{\overleftarrow{\rho}}{2} \|\bar{x} - \bar{n}\|^2 \right] \quad (20)$$

$$\text{s.t. } \|(\alpha \bar{x} + (1 - \alpha) \underline{x}) - (\beta x_R + (1 - \beta) x_L)\| \geq r \text{ for all } \alpha, \beta \in [0, 1]. \quad (21)$$

Observe that the term $\frac{\overleftarrow{\rho}}{2} \|\underline{x} - \underline{n}\|^2$ equals the energy of a spring with zero rest-length and elastic coefficient $\overleftarrow{\rho}$ whose end points are at positions \underline{x} and \underline{n} . The same interpretation applies for the

second term in (20). With this interpretation in mind, the non-convex constraint (21) means that the line from \underline{x} to \bar{x} cannot cross the region swept out by a circle of radius r that moves from x_L to x_R . We call this region \mathcal{R} . Figure 6-left shows a feasible solution and an unfeasible solution under this interpretation. When the line from \underline{n} to \bar{n} does not cross \mathcal{R} , the solution of (20)-(21) is



Figure 6: Left: Feasible solution (blue) and unfeasible solution (red). Right: Two different feasible configurations of the springs-slab system. Each represented in different color.

$\underline{x} = \underline{n}$ and $\bar{x} = \bar{n}$. In general however, \underline{x} and \bar{x} adopt the minimum energy configuration of a system with two zero rest-length springs, with end points $(\underline{n}, \underline{x})$ and (\bar{n}, \bar{x}) and elastic coefficients $\overleftarrow{\rho}$ and $\overleftarrow{\bar{\rho}}$, and with a hard extensible slab, connecting \underline{x} to \bar{x} , that cannot go over region \mathcal{R} . The slab can be extended without spending any energy. Figure 6-right shows two feasible configurations of the system of springs and slab when $\underline{x} = \underline{n}$ and $\bar{x} = \bar{n}$ cannot be a feasible solution.

It is possible that this minimizer receives two zero-weight messages from its neighboring equality nodes, i.e., $\overleftarrow{\rho} = \overleftarrow{\bar{\rho}} = 0$. This would correspond to not having any spring connecting point \underline{n} to \underline{x} and \bar{n} to \bar{x} . The mechanic system would then be indeterminate. When this is the case, we solve the mechanic system in the limit when $\overleftarrow{\rho} = \overleftarrow{\bar{\rho}} \rightarrow 0^+$. In terms of software implementation, this is achieved by replacing $\overleftarrow{\rho}$ and $\overleftarrow{\bar{\rho}}$ by small equal values.

In Section I we explain how to compute the minimum energy configuration of this system quickly. In other words, we show that the minimizer g^{wall} can be implemented efficiently.

E Energy minimizer

The energy minimizer solves the quadratic optimization problem

$$\min_{\{\underline{x}, \bar{x}\}} \left[C \|\underline{x} - \bar{x}\|^2 + (\overleftarrow{\rho}/2) \|\underline{x} - \underline{n}\|^2 + (\overleftarrow{\bar{\rho}}/2) \|\bar{x} - \bar{n}\|^2 \right].$$

From the first order optimality conditions we get $2C(\underline{x} - \bar{x}) + \overleftarrow{\rho}(\underline{x} - \underline{n}) = 0$ and $2C(\bar{x} - \underline{x}) + \overleftarrow{\bar{\rho}}(\bar{x} - \bar{n}) = 0$. Solving for \underline{x} and \bar{x} we obtain,

$$\underline{x} = \frac{\overleftarrow{\rho} \overleftarrow{\bar{\rho}} \underline{n} + 2C(\overleftarrow{\rho} \bar{n} + \overleftarrow{\bar{\rho}} \underline{n})}{2C(\overleftarrow{\rho} + \overleftarrow{\bar{\rho}}) + \overleftarrow{\rho} \overleftarrow{\bar{\rho}}}, \quad \bar{x} = \frac{\overleftarrow{\rho} \overleftarrow{\bar{\rho}} \bar{n} + 2C(\overleftarrow{\rho} \underline{n} + \overleftarrow{\bar{\rho}} \bar{n})}{2C(\overleftarrow{\rho} + \overleftarrow{\bar{\rho}}) + \overleftarrow{\rho} \overleftarrow{\bar{\rho}}}. \quad (22)$$

If the energy minimizer receives $\overleftarrow{\rho} = \overleftarrow{\bar{\rho}} = 0$, we resolve the indeterminacy in computing (22) by letting $\overleftarrow{\rho} = \overleftarrow{\bar{\rho}} \rightarrow 0^+$.

F Maximum velocity minimizer

This minimizer solves the convex problem $\text{minimize}_{\{\underline{x}, \bar{x}\}} \left[(\overleftarrow{\rho}/2) \|\underline{x} - \underline{n}\|^2 + (\overleftarrow{\bar{\rho}}/2) \|\bar{x} - \bar{n}\|^2 \right]$ subject to $\|\underline{x} - \bar{x}\| \leq C$. If $\|\underline{n} - \bar{n}\| \leq C$ then $\underline{x} = \underline{n}$ and $\bar{x} = \bar{n}$. Otherwise, the constraint is active, and, using the KKT conditions, we have $\overleftarrow{\rho}(\underline{x} - \underline{n}) = -\lambda(\underline{x} - \bar{x})$ and $\overleftarrow{\bar{\rho}}(\bar{x} - \bar{n}) = -\lambda(\bar{x} - \underline{x})$ where $\lambda \neq 0$ is such that $\|\underline{x} - \bar{x}\| = C$. Solving for \underline{x} and \bar{x} we get,

$$\underline{x} = \frac{\overleftarrow{\rho} (\overleftarrow{\bar{\rho}} + \lambda) \underline{n} + \lambda \overleftarrow{\bar{\rho}} \bar{n}}{\overleftarrow{\rho} \overleftarrow{\bar{\rho}} + \lambda(\overleftarrow{\rho} + \overleftarrow{\bar{\rho}})}, \quad \bar{x} = \frac{\overleftarrow{\bar{\rho}} (\overleftarrow{\rho} + \lambda) \bar{n} + \lambda \overleftarrow{\rho} \underline{n}}{\overleftarrow{\rho} \overleftarrow{\bar{\rho}} + \lambda(\overleftarrow{\rho} + \overleftarrow{\bar{\rho}})}. \quad (23)$$

To find the solution we just need to determine λ . Computing the difference between the above expressions we get,

$$\underline{x} - \bar{x} = \frac{\underline{n} - \bar{n}}{1 + \left(\frac{1}{\underline{\rho}} + \frac{1}{\bar{\rho}}\right)\lambda}. \quad (24)$$

Taking the norm of the right hand side and setting it equal to C we get

$$\lambda = \pm \frac{(\|\underline{n} - \bar{n}\|/C) - 1}{\underline{\rho}^{-1} + \bar{\rho}^{-1}}. \quad (25)$$

Now examine equation (24). Starting from an $\underline{n} - \bar{n}$ such that $\|\underline{n} - \bar{n}\| > C$, the fastest way to get to $\underline{x} - \bar{x}$ with $\|\underline{x} - \bar{x}\| = C$ is to increase $\lambda > 0$. Hence, in (25), we should choose the positive solution, i.e.

$$\lambda = \frac{(\|\underline{n} - \bar{n}\|/C) - 1}{\underline{\rho}^{-1} + \bar{\rho}^{-1}}. \quad (26)$$

If the maximum velocity minimizer receives $\underline{\rho} = \bar{\rho} = 0$, we resolve any indeterminacy by letting $\underline{\rho} = \bar{\rho} \rightarrow 0^+$. In software, this is achieved by setting $\underline{\rho}$ equal to $\bar{\rho}$ equal to some small value.

G Minimum velocity minimizer

This minimizer can be computed in a very similar way to the maximum velocity minimizer. If $\|\underline{n} - \bar{n}\| \geq C$, then $\underline{x} = \underline{n}$ and $\bar{x} = \bar{n}$. Otherwise, from the KKT conditions, we again obtain equation (23). The difference $\underline{x} - \bar{x}$ is again the expression (24). Now, however, starting from $\underline{n} - \bar{n}$ such that $\|\underline{n} - \bar{n}\| > C$, the fastest way to get to $\underline{x} - \bar{x}$ with $\|\underline{x} - \bar{x}\| = C$, is to decrease $\lambda < 0$. Hence, in (25), we should choose the negative solution, i.e. (26) holds again. If the minimum velocity minimizer receives $\underline{\rho} = \bar{\rho} = 0$, we resolve any indeterminacy by letting $\underline{\rho} = \bar{\rho} \rightarrow 0^+$. In software, this is achieved by setting $\underline{\rho}$ equal to $\bar{\rho}$ equal to some small value.

H Velocity obstacle minimizers

In this section we explain how to write the minimizers associated to each of the terms in equation (13) using the minimizers g^{coll} , g^{wall} and g^{cost} for global planning.

First however, we briefly describe the bipartite graph that connects all these minimizers together. The bipartite graph for this problem has a $g^{\text{VO coll}}$ minimizer-node connecting every pair of equality-nodes. There is one equality-node per variable in $\{x_i\}$. Recall that each of these variables describes the position of an agent at the end of a planning epoch. Each equality-node is also connected to a separate $g^{\text{VO cost}}$ minimizer-node. Finally, for obstacles in 2D, every equality node is also connected to several $g^{\text{VO wall}}$ minimizer-nodes, one per obstacle.

We start by describing the minimizer associated to the terms $\{f_{r_i, r_j}^{\text{coll}}(x_i(0), x_i, x_j(0), x_j)\}$ in equation (13). This is given by

$$g^{\text{VO coll}}(\underline{n}, \underline{n}', \underline{\rho}, \underline{\rho}', x^0, x'^0, r, r') = \arg \min_{\{\underline{x}, \underline{x}'\}} \left[f_{r, r'}^{\text{coll}}(x^0, \underline{x}, x'^0, \underline{x}') + \frac{\underline{\rho}}{2} \|\underline{x} - \underline{n}\|^2 + \frac{\underline{\rho}'}{2} \|\underline{x}' - \underline{n}'\|^2 \right].$$

The messages \underline{n} and \underline{n}' , and corresponding certainty weights $\underline{\rho}$ and $\underline{\rho}'$, come from the equality-nodes associated to the end position of two agents of radius r and r' that, during one time epoch, move from their initial positions x^0 and x'^0 to \underline{x} and \underline{x}' without colliding.

The outgoing weights $\vec{\rho}$ and $\vec{\rho}'$ associated to the variables \underline{x} and \underline{x}' are determined in the following way. If an agent of radius r moving from x^0 to \underline{n} does not collide with an agent of radius r' moving from x'^0 to \underline{n}' , the minimizer will not propose a new trajectory for them, i.e., the minimizer will return $\underline{x} = \underline{n}$ and $\underline{x}' = \underline{n}'$. Hence, in this case, we set all outgoing weights equal to 0, signaling to neighboring equality-nodes that the minimizer wants to have no say when try to reach consensus. Otherwise, we set all outgoing weights equal to ρ_0 .

For this minimizer, by direct substitution one sees that,

$$g^{\text{VO coll}}(\underline{n}, \underline{n}', \overleftarrow{\underline{\rho}}, \overleftarrow{\underline{\rho}'}, x^0, x'^0, r, r') = g^{\text{coll}}(x^0, \underline{n}, x'^0, \underline{n}', +\infty, \overleftarrow{\underline{\rho}}, +\infty, \overleftarrow{\underline{\rho}'}, r, r'). \quad (27)$$

Above we are using a notation where, given a function f , $f(+\infty) \equiv \lim_{x \rightarrow +\infty} f(x)$. In software, this is implemented by replacing $+\infty$ by a very large value.

In a very similar way, the minimizer associated to the terms $\{f_{x_{Rk}, x_{Lk}, r_i}^{\text{wall}}(x_i(0), x_i)\}$ can be written using the agent-obstacle minimizer for the global planning problem. Concretely,

$$g^{\text{VO wall}}(\underline{n}, x_L, x_R, r, \overleftarrow{\underline{\rho}}) = g^{\text{wall}}(x^0, \underline{n}, r, x_L, x_R, +\infty, \overleftarrow{\underline{\rho}}). \quad (28)$$

For this minimizer, the rule to set the outgoing weights is the following. If an agent of radius r can move from x^0 to \underline{n} without colliding with the line segment $[x_L x_R]$ then set all outgoing weights to 0. Otherwise set them to ρ_0 .

Finally, we turn to the the minimizer associated to the terms $\{f_{C_i}^{\text{cost}}(x_i, x_i^{\text{ref}})\}$. This minimizer receives as input a message \underline{n} , with corresponding certainty weight $\overleftarrow{\underline{\rho}}$, from the equality-minimizer associated to the position of an agent at the end of a time epoch and outputs a local estimate, \underline{x} , of its position at the end of the epoch. It also receives as parameter a reference position x^{ref} and a cost c of \underline{x} deviating from it. To be concrete, its output is chosen uniformly at random from following arg min set,

$$g^{\text{VO cost}}(\underline{n}, \overleftarrow{\underline{\rho}}, x^{\text{ref}}, c) = \arg \min_{\underline{x}} \left[f_c^{\text{cost}}(\underline{x}, x^{\text{ref}}) + \frac{\overleftarrow{\underline{\rho}}}{2} \|\underline{x} - \underline{n}\|^2 \right] \quad (29)$$

$$= \arg \min_{\underline{x}} \left[c \|\underline{x} - x^{\text{ref}}\|^2 + \frac{\overleftarrow{\underline{\rho}}}{2} \|\underline{x} - \underline{n}\|^2 \right]. \quad (30)$$

The outgoing weights for this minimizer are always set to ρ_0 .

Again by direct substitution we see that,

$$g^{\text{VO cost}}(\underline{n}, \overleftarrow{\underline{\rho}}, x^{\text{ref}}, c) = g^{\text{cost}}(x^{\text{ref}}, \underline{n}, +\infty, \overleftarrow{\underline{\rho}}, c), \quad (31)$$

where in g^{cost} we are using the energy minimizer for the global planning problem.

I Mechanical analog

In this section we explain how to compute the minimum energy configuration of the springs-slab system described in Section D. Basically, it reduces to computing the minimum of a one-dimensional real function over a closed interval.

Given \underline{n} and \bar{n} , two main scenarios need to be considered.

1. If $\underline{n}, \bar{n} \notin \mathcal{R}$ and $[\underline{n} \bar{n}] \cap \mathcal{R} = \emptyset$, i.e. the segment from \underline{n} to \bar{n} does not intersect \mathcal{R} , then $\underline{x} = \underline{n}$ and $\bar{x} = \bar{n}$.
2. Otherwise, because there might be multiple local minima, i.e. multiple stable static configurations, we need to compare the energy of the following two configurations and return the one with minimum energy.
 - (a) The slab is tangent to \mathcal{R} , for example as in the blue configuration of Figure 6-right.
 - (b) One of the springs is fully compressed and exactly one end of the slab is touching the boundary of \mathcal{R} , for example as in the red configuration of Figure 6-right.

Let us compute the energy for Scenario 2a. The slab can be tangent to \mathcal{R} in many different ways. However, the arrangement must always satisfy two properties. First, the point of contact, p , between the slab and \mathcal{R} touches either the boundary of the semi-circle centered at x_L or the boundary of the semi-circle centered at x_R . Second, because extending/compressing the slab costs zero energy, the slab must be orthogonal to the line segment $[\underline{n} \underline{x}]$ and to the line segment $[\bar{n} \bar{x}]$.

The first observation allows us to express p using the map $P(\theta) : [0, 2\pi] \mapsto \mathbb{R}^2$ between the direction of the slab and the point of contact at boundary of the semi-circles,

$$P(\theta) = \begin{cases} x_R + r\hat{n}(\theta) & \langle x_R - x_L, \hat{n}(\theta) \rangle \geq \langle x_L - x_R, x_R \rangle \\ x_L + r\hat{n}(\theta) & \text{otherwise} \end{cases} \quad (32)$$

where $\hat{n}(\theta) = \{\cos(\theta), \sin(\theta)\}$. Specifically, there is a $\theta^0 \in [0, 2\pi]$ such that $p = P(\theta^0)$. The second observation tells us that $\underline{x} = \underline{n} + \underline{\gamma}\hat{n}(\theta^0)$ and $\bar{x} = \bar{n} + \bar{\gamma}\hat{n}(\theta^0)$ where $\underline{\gamma}$ and $\bar{\gamma}$ can be determined using the orthogonality conditions,

$$\langle \underline{x} - P(\theta^0), \hat{n}(\theta^0) \rangle = 0 \Rightarrow \underline{\gamma} = \langle P(\theta^0) - \underline{n}, \hat{n}(\theta^0) \rangle \quad (33)$$

$$\langle \bar{x} - P(\theta^0), \hat{n}(\theta^0) \rangle = 0 \Rightarrow \bar{\gamma} = \langle P(\theta^0) - \bar{n}, \hat{n}(\theta^0) \rangle. \quad (34)$$

Therefore, the minimum energy configuration over all tangent configurations, which is fully determined by θ^0 , must satisfy

$$\theta^0 \in \arg \min_{\theta \in [0, 2\pi]} E_{\text{tangent}}(\theta) \quad \text{where,} \quad (35)$$

$$E_{\text{tangent}}(\theta) = \frac{\underline{\rho}}{2} (\langle P(\theta) - \underline{n}, \hat{n}(\theta) \rangle)^2 + \frac{\bar{\rho}}{2} (\langle P(\theta) - \bar{n}, \hat{n}(\theta) \rangle)^2. \quad (36)$$

Problem (35) involves minimizing the one-dimensional function (36). Figure 7-left shows the typical behavior of $E_{\text{tangent}}(\theta)$. It is non-differentiable in at most 2 points. When the agent-obstacle




Figure 7: Typical behavior of $E_{\text{tangent}}(\theta)$ for the agent-obstacle minimizer (left) and for the agent-agent minimizer (right).

minimizer is used to solve the agent-agent minimizer, the function becomes smooth and has second derivative throughout all its domain, see Figure 7-right. In the numerical results of Section 4, our implementation of the agent-agent minimizer uses Newton's method to solve (35). To find the global minimum, we apply Newton's method starting from four equally-space points in $[0, 2\pi]$. To produce the video accompanying this appendix, our implementation of the agent-obstacle minimizer solves (35) by scanning points in $[0, 2\pi]$ with a step size of $2\pi/1000$. In this case, it is obvious there is room for improvement in speed and accuracy by choosing smarter ways in which to solve (35).

To compute the energy for Scenario 2b, we need to determine which of the springs is fully contracted, or which side of the slab is touching \mathcal{R} . If $\underline{n} \in \mathcal{R}$ and $\bar{n} \notin \mathcal{R}$ then $\bar{x} = \bar{n}$ and \underline{x} is the point in the boundary of \mathcal{R} closest to \underline{n} such that $[\underline{x} \bar{x}]$ does not intersect \mathcal{R} . Since the boundary of \mathcal{R} is formed by parts of the boundary of two circles and of two lines, this closest point can be computed in closed form. If $\bar{n} \in \mathcal{R}$ and $\underline{n} \notin \mathcal{R}$ the situation is the opposite. If $\underline{n}, \bar{n} \in \mathcal{R}$, then we know we cannot be in Scenario 2b. Finally, if $\underline{n}, \bar{n} \notin \mathcal{R}$, we compute the energy assuming that $\underline{x} = \underline{n}$ and then assuming that $\bar{x} = \bar{n}$ and take the configuration with smallest energy between them.

J Comment on our algorithm

Note that our algorithm does *not* possess anytime guarantees and, if stopped earlier, the trajectories might have collisions. However, if stopped early, a suboptimal set of non-colliding trajectories can be found at very low computational cost by using our algorithm to solve the feasibility problem in (2) starting from the state of the algorithm at stop time. In addition, although dynamic/static obstacles

can be seamlessly integrated into our framework, a solution must be recomputed (as is the case with A* or RRT*) if their trajectories/positions change unexpectedly. This being said, in our algorithm, if a new piece of information is received, the previous solution can be used as the initial guess, potentially decreasing convergence time. Note that in some scenarios, a low-cost local-planning approach, such as the one presented in Section 5, can be beneficial.

K Comment on the scaling of convergence time with p

Based on [13], we think that in non-adversarial configurations, in contrast to CONF1 where deadlocks are likely, the scaling of convergence time with p is not exponential. Our reasoning is based on a connection between trajectory planning and disk packing. For example, minimum-energy trajectory planning using piece-wise linear trajectories is related to, although not the same as, packing disks in multiple 2D layers, where the two matched disks between consecutive layers generate a larger cost when far away from each other. The numerical results in [13] report that, for disk packing, the runtime of ADMM and TWA is no more than polynomial in the number of disks and we believe the runtime for trajectory planning for non-adversarial configurations has a similar complexity. We interpret the seemingly exponential curve of convergence time versus p for $n = 8$ in Figure 2-left as an atypical, adversarial scenario. By comparison, in Figure 3-left, which assumes randomly chosen initial and final points and also minimization of energy, the dashed-blue curve of runtime versus p for $n = 8$ does not appear to exhibit exponential growth.