Human-Inspired Distributed Collision Avoidance

Ross A. Knepper and Daniela Rus

***Abstract*—We present a distributed collision avoidance algorithm for multiple moving robots that is model-predictive and sampling-based. Unlike purely reactive approaches, the proposed algorithm incorporates arbitrary trajectories as generated by a motion planner running on the navigating robot and predicted trajectories of its co-agents (robots or humans). Our approach, inspired by pedestrian navigation in crowded spaces, draws from the ethnography literature on pedestrian interaction. We propose a simple two-phase algorithm in which agents first cooperate to avoid each other and then engage in *civil inattention*, just as pedestrians look away as they pass by each other. This process entails a pedestrian bargain in which all agents act competently to avoid each other and, once resolution is achieved, to avoid interfering with others' planned trajectories.**

Sampling-based planners have often been employed to enable robots to avoid static and even dynamic obstacles. If the obstacle is an intelligent agent, such as a human or another robot, this problem is complicated by the difficulty in predicting the agent's reaction to the robot's own movements. Multi-robot collision avoidance algorithms have been proposed that can incorporate the reactions of fellow robots. However, many of these, such as reciprocal velocity obstacles [1], are purely reactive and therefore cannot incorporate general predictions about an agent's intended trajectory.

We describe a distributed, sampling-based, cooperative collision-avoidance algorithm that is predictive, reactive, and reciprocal, thus leaving the system of robots free from instability. Our key insight—stemming from observations of humans navigating through pedestrian traffic—is that people perform pairwise mutual avoidance by reacting to each interfering person at most once. Robot negotiations are patterned after the **Pedestrian Avoidance Procedure:**

1) The interaction begins when a person perceives a possible future collision with another pedestrian.
2) *If the pedestrian appears competent and engaged,* then the person makes a visible move to correct their trajectory by about half of the amount required to fully avoid the collision.
3) *Otherwise,* the person makes a full effort to avoid the pedestrian.
4) Finally, the person resolves the interaction, often by initiating *civil inattention*.

Civil inattention describes the act of deliberately looking away in order to signal the belief that this particular interaction is resolved and consequently that one will not react again. Such refusal to react is an important aspect of the algorithm, as it prevents oscillatory behavior in nearly all cases.

The Pedestrian Avoidance Procedure can be easily recreated within the sampling-based planning paradigm (Fig. 1). In this framework, a robot i tracks other agents in its neighborhood (whether robot or human). For an agent j, it maintains a tuple (t_j, r_{ij}, r_{ji}), in which t_j stores a prediction of agent j's trajectory, and r_{ab} is a boolean indicator of whether agent a has recently reacted to agent b.

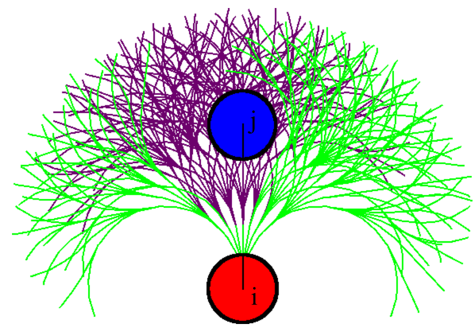

Fig. 1. During the interaction phase, robot i rejects paths that result in an overlap $> 50\%$ with the predicted trajectory of agent j (a robot or human), who is expected to reciprocate. A right-side sampling bias follows social convention. After the interaction ends, each agent treats the other as a non-reactive moving obstacle, requiring 100% avoidance.

When the robot predicts a collision with an agent's trajectory, the robot conditions its sampling strategy on whether or not that agent has already finished reacting to the robot. During an interaction, the robot anticipates that the agent will cooperatively avoid it, and so the robot rejects samples that would result in an overlap with the agent's trajectory by 50% or more—that is, less than half overlap is permitted. Once agent j sets r_{ji} to true, robot i instead rejects samples that intersect agent j's trajectory at all, since the agent has promised not to react further. The pedestrian-inspired collision avoidance algorithm provides several important properties:

- The algorithm is immune to deadlock in environments with little to no static obstacle clutter. Note that a different type of reasoning is required to solve maze-like dense clutter problems.
- The algorithm functions correctly and safely under the assumptions that all agents are competent, cooperative, and aware of each other.

We performed scalability experiments to assess multirobot performance. We demonstrated the algorithm on up to four KUKA youBot holonomic bases. A separate instance of the algorithm ran on each robot with communications via ROS over WiFi. In simulation, we demonstrated up to sixteen youBots in Gazebo, with separate motion planning processes running on a single dual-core CPU. On a position-exchange task, we observed a runtime to scale linearly with the number of interacting robots.

Although we do not explicitly address robot motion planning in the presence of humans, the sampling-based motion planning algorithm we present is simultaneously predictive and reactive, thus holding promise for integration into human environments. Furthermore, it is fully distributed and scales efficiently to large groups.

References

[1] J. van den Berg, M. Lin, and D. Manocha. Reciprocal Velocity Obstacles for Real-Time Multi-Agent Navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2008.