The Colliding Reciprocal Dance Problem: A Mitigation Strategy with Application to Automotive Active Safety Systems

Jeffrey Kane Johnson[1]

***Abstract*— A *reciprocal dance* occurs when two mobile agents attempt to pass each other but incompatible interaction models result in repeated attempts to take mutually blocking actions. Such a situation often results in deadlock, but in systems with significant inertial constraints, it can result in collision. This paper presents this colliding variant of the reciprocal dance, how it arises, and a mitigation strategy that can improve safety without sacrificing flexibility. A demonstration is provided in the context of automotive active safety.**

Supplemental Materials

Data sets and a demonstration video are available at: https://jeffreykanejohnson.com/publications/#crd.

I. Introduction

A familiar problem for people walking through crowded areas is the a phenomenon where two people are unable to pass each other even if sufficient space for passage exists. This happens when incompatible interaction models result in attempts to take mutually blocking actions. The repeated selection of incompatible actions results in an oscillating deadlock that we refer to as a "Reciprocal Dance."

For pedestrians, a reciprocal dance is typically just a nuisance. But for agents with inertial constraints, a reciprocal dance can lead to collision with severe, or even fatal, consequences. We will show that a simple mitigation strategy can help reduce the severity of dangerous reciprocal dance situations or avoid them entirely.

We choose to mitigate rather than solve the problem for the simple reason that an exact solution is not realistically feasible. As the next section discusses, any exact solution will be intractable and therefore not computable. Following that, the paper will address the colliding variant of the reciprocal dance, how it arises, our mitigation strategy, and two examples in the context of automotive active safety.

II. Background

The problem of motion planning for multiple agents following non-fixed trajectories is generally framed in terms of sequential decision making. These problems can often be efficiently solved in centralized cases through formulation as a type of Markov decision process (MDP), which belongs to complexity class P [1]. Decentralized problems present a different set of challenges. It was shown in [2] that for cooperative agents (i.e., agents that share a reward function) this class of problems is at least complete for NEXP in both

[1]Jeff is Principal at Mapless AI, Inc., `jeff@mapless.ai`

(a) Bird's eye view (b) Front camera view

Fig. 1: The vehicle (left pane, bottom) drives into stationary traffic, relying on the safety system to modulate throttle.

(a) Bird's eye view (b) Front camera view

Fig. 2: The vehicle (left pane, between white and blue vehicles) navigates moving traffic, relying on the safety system to modulate steering and throttle.

the jointly fully observable (DEC-MDP) and jointly partially observable (DEC-POMDP) cases. It was shown in [3] that the partially observable stochastic game (POSG), which is the non-cooperative version of this problem (i.e., the problem in which agents do not share a reward function), is complete for NEXP^{NP}. There is little chance of computing exact solutions for anything but toy problems in these complexity classes [4], which is unfortunate because formulations like the POSG provide realistic models of many real-world scenarios, such as vehicle traffic.

Thus, problem complexity makes dealing with multi-agent systems in an exact way almost impossible. But it gets worse: even approximations are often intractable if any guarantee is required of how close to correct the approximation is [5]. This effectively forces the use of loose approximations or heuristics in order to compute solutions, which is why we have chosen to pursue a mitigation strategy rather than an exact solution. A metaheuristic useful for deriving such approximate solutions comes from decentralized control theory in the form of *person-by-person* solution techniques [6], [7]. In these, each agent assumes arbitrary control strategies for all other agents and optimizes its own actions according to the assumed strategies. These techniques provide no global optimality guarantees, but they can be made safe and computationally efficient. Dynamic Virtual Bumper [8], Dynamic Window [9], and Velocity Obstacle (VO) [10] approaches, employ this strategy to great success for low-order systems. For higher-order systems, the VO representation has been modified and extended [11], but it becomes more difficult to account for higher-order constraints in a fundamentally first-order representation. The Selective Determinism [12] framework used in this paper is person-by-person approach similar to the Dynamic Window, but is designed for multi-agent systems with arbitrary dynamics that allows safety guarantees to be maintained.

III. Formalizing and Mitigating the Colliding Reciprocal Dance Problem

A. Selective Determinism

In many multi-agent navigation problems, it can be assumed that agents are non-adversarial and will prefer self preservation to goal-directed progress. This assumption can be exploited to safely factor behavioral interaction effects out of sequential decision and enable tractable solutions. *Selective Determinism* [12] is a framework that formalizes this process. The following definitions will be used to derive our mitigation strategy using the framework.

Definition 1. The *agent state* $A(\mathbf{x})$ is a state representation containing both the dynamic state of an agent A and the volume of workspace it occupies at state $\mathbf{x}$.

Definition 2. A *stopping path* is the minimum space an agent needs to occupy in order to come to a stop[1] along a given path. Also referred to as a *contingency plan*. See Figure 3a.

Definition 3. A *stopping region* is the disjoint union of all stopping paths defined over the set of all followable paths. Also referred to as a *contingency set*. See Figure 3b.

Property 1. *Stopping path disjointness* is the property that a system is guaranteed the ability to remain collision free if all agents have one or more stopping paths that are disjoint from the stopping regions of all other agents.

Property 1 guarantees the existence of a collision-free contingency plan for all agents, which guarantees that collision can always be avoided. Selective determinism provides a mechanism for all agents to maintain this property independently, provided that all agents invoke a contingency if *not* doing so will lead to the property being violated. As the name of Property 1 implies, the mechanism involves the computation of stopping regions and stopping paths.

By definition, stopping regions depend only on agent dynamics and not interactions or behaviors. Thus, agents can compute the regions for themselves and others without regard to the partial observability of decision processes. This makes it possible for each agent to *tractably* compute actions that ensure it maintains at least one collision free stopping path (i.e., contingency plan), which guarantees that collisions can always be efficiently avoided.

[1] In this context "stop" could mean zero relative, or zero absolute, velocity.

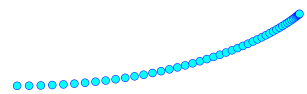

(a) Illustration of a *stopping path* for a disc agent along a curved path moving from left to right.

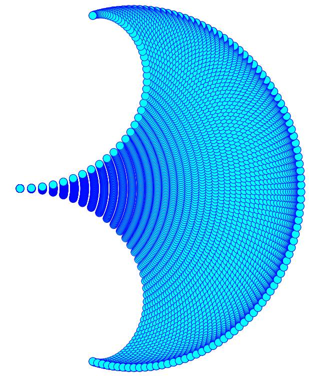

(b) Illustration of the *stopping region* for the disc agent from (a).

Fig. 3: Illustrations of a stopping path and stopping region for a moving 2D disc agent with unicycle dynamics following constant control trajectories.

It is this, rather than policy optimality, that Selective Determinism uses to ensure safety. Because safety is not dependent on policy optimality, agents can also trade optimal interaction models for computationally simpler ones. Though enabling from a tractability standpoint, this trade-off also allows certain pathological situations.

B. Reciprocal Dance as Pathological Selective Determinism

In many applications, interactions tend to be simple and incompatibilities transient, which means agents with simplified models generally navigate successfully. But what if interaction models are *consistently* incompatible? This pathological case can provide a revised definition of a reciprocal dance:

Definition 4. A *Reciprocal Dance* is a situation in Selective Determinism when mutually incompatible interaction models cause a deadlock of repeated contingency invocation.

C. The Colliding Reciprocal Dance

A reciprocal dance can be *colliding* if it occurs in a system where Property 1 is not maintained. Unfortunately, maintaining this property can sometimes be non-trivial. In order for it to hold, the boundary of a *computed* stopping region must be conservative with respect to the boundary of the *true* stopping region. In systems with significant inertial constraints this can be challenging to guarantee because the stopping paths sweep out large and complicated volumes that can be difficult to usefully approximate. However, a simple mitigation strategy can help significantly.

D. A Mitigation Strategy: Constraint Tightening

The following definitions will be useful for outlining the mitigation strategy.

Definition 5. A *constraint set* is a set of position-indexed constraint functions that define dynamic constraints for feasible motion at each position along the path for a given agent.

Definition 6. A *nominal constraint set* is a constraint set that defines all feasible motion of an agent.

Definition 7. A *contingency constraint set* is a constraint set that defines stopping paths.

Definition 8. A *constraint set collection* is a path-indexed set of nominal and contingency constraint sets.

As a mitigation strategy, we propose to bias an agent's dynamic constraints to those of its current stopping region. By exploiting the definition of a stopping region, this is straightforward to do. The agent has some set of paths available to it as well as a constraint set collection that defines nominal and contingency constraint sets for those paths. The stopping region has the same set of paths coupled with *contingency* constraint sets that most quickly bring the agent to a stop. *Constraint tightening* is the process of adjusting the bounds of the nominal constraints to be nearer those of the contingency constraints. This strategy is effectively a type of adaptive damping that can reduce the risk of violating Property 1 by reducing the likelihood that stopping paths are not disjoint.

For the intuition behind the strategy, consider the constraints on path position derivatives, which govern how an agent moves[2] along a path. Lowering the bound on any of these derivatives reduces the amount of kinetic energy an agent can attain, which reduces the displacement along the path necessary to stop. For example, consider a point agent moving along a frictionless curve with the below state function. The function describes the first time derivative of arc length position $\dot{s}$ along the curve with respect to a higher-order time derivative $s^{(n)}$ and time t:

$$\dot{s}\left(s^{(n)}, t\right) = \sum_{k=0}^{n-3} a_k t^k + \frac{1}{(n-1)!} s^{(n)} t^{n-1}$$

This is a simple n^{th}-order integrated control system similar to the kind used in Frenét-frame motion planning [13]. For this state function it is clear that for $s_1^{(n)} \leq s_2^{(n)}$:

$$\dot{s}_1\left(s_1^{(n)}, t\right) \leq \dot{s}_2\left(s_2^{(n)}, t\right)$$

Thus, assuming $\dot{s}_1, \dot{s}_2 \geq 0$, if the same deceleration is applied to both, the arc length displacement needed to reach zero speed from $\dot{s}_1$ will be no more than that from $\dot{s}_2$, which makes the relationship between stopping path length and value of $s^{(n)}$ order preserving, or monotonic. This is an intuitive idea that we codify as an assumption for systems using our mitigation strategy:

Assumption 1. *The workspace volume needed to bring a moving system to rest is monotonically non-decreasing with respect to how quickly the system is capable of changing its dynamic state.*

We make Assumption 1 explicit to avoid the need to define the mitigation strategy with respect to a specific motion model or deal with pathological cases where it neither describes nor reasonably approximates a system.

Lemma 1 now derives the relationship between stopping region volume and constraint tightening.

Lemma 1. *For path set $\mathcal{P}$ and nominal and contingency constraint set collections $\mathcal{C}^n$ and $\mathcal{C}^c$, let $\mathcal{C}^\star$ be a tightened constraint set collection. For any path $P_i \in \mathcal{P}$ and agent state $A(\mathbf{x})$, let $u^\star$ be a control generated with respect to $\mathcal{C}^\star$ and u^n a control generated with respect to $\mathcal{C}^n$. Let SR^n be the stopping region that results from applying u^n to $A(\mathbf{x})$ and let $SR^\star$ be the stopping region for $u^\star$. If $\mathcal{W}(\cdot)$ is function that computes workspace volumes of stopping regions, then $\mathcal{W}(SR^\star) \leq \mathcal{W}(SR^n)$.*

Proof. Trivially, if $u^n = u^\star$, then $\mathcal{W}(SR^\star) = \mathcal{W}(SR^n)$. Otherwise, by definition of constraint tightening, the magnitude of state change due to $u^\star$ will not be greater than that due to u^n. Under Assumption 1, it follows from Definition 3 that $\mathcal{W}(SR^\star) \leq \mathcal{W}(SR^n)$. □

Ideally, we would like the workspace volume to be *smaller* than it otherwise would have been. Unfortunately, this can't be guaranteed in general because controls can saturate (e.g. $\dot{s}_1 = \dot{s}_2$ for $u_1 > u_2$) or the stopping regions can have arbitrary self-intersections. For example, a sphere spinning in place would have a stopping region workspace volume exactly equal to its own volume, so it would be impossible for the volume of the region to ever be smaller. In many applications, however, constraint bounds are more strongly correlated to workspace stopping volume. Consider the stopping region shown in Figure 3b: any reduction in stopping path length directly correlates to smaller workspace volume for the stopping region.

Constraint tightening can slow, or even reverse, the growth of the true stopping region volume proportional to the degree of tightening. In practice, it is quite useful to set the degree of tightening adaptively as some function of proximity. In this work, we define proximity as the minimum time t^c before a contingency may need to invoked. We find a temporal proximity simpler to implement and to lead to more consistent and intuitive behavior.

We perform tightening by computing a scaling factor $\gamma \in [0, 1]$ that scales from the contingency bound at 0 to the nominal bound at 1:

$$c^\star = (1 - \gamma)c^c + \gamma c^n$$

Note that this scaling relies on another assumption:

[2]For simplicity of discussion we assume that agent orientation is determined by path position, but in general this need not be the case.

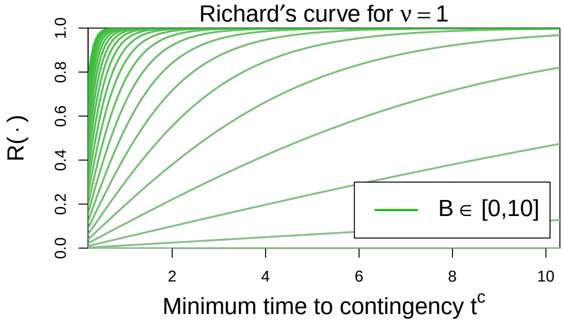

(a) The plot shows R plotted for sampled values of B, where the plot shifts from the x-axis (more conservative) to the y-axis (less conservative) as B approaches 10.

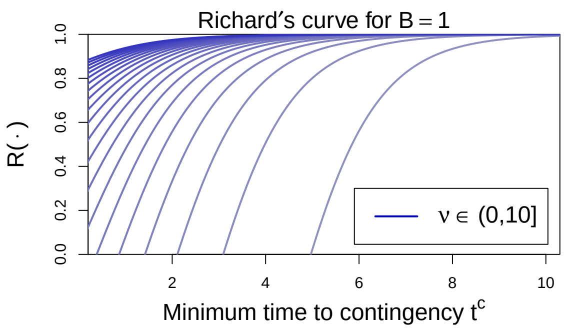

(b) The plot shows R plotted for sampled values of ν, where the plot moves from left to right as ν approaches 10. For $\nu = 1$ the plot crosses the origin.

Fig. 4: The effects of adjusting parameter values B and ν.

Assumption 2. *A safe and feasible control can be computed with respect to every $c^\star \in [c^c, c^n]$.*

Maintaining and verifying Assumption 2 is highly model dependent and can be tricky if feasible constraint regions are not simply connected or if contingency sets are not contained within nominal sets. Future work will investigate formal methods for handling Assumption 2.

The scaling factor γ is computed in part using a generalized logistic function, or Richard's curve, R:

$$\gamma\left(t^c\right) = \begin{cases} 0 & t^c \leq 0 \\ \max\left(0, R\left(t^c\right)\right) & t^c > 0 \end{cases}$$

We partially specify the parameters of R such that it maps to $[-1, 1]$:

$$R\left(t^c\right) = \frac{2}{\left(1 + e^{-Bt^c}\right)^{1/\nu}} - 1$$

This leaves parameters $B \in [0, \infty)$ and $\nu \in (0, \infty)$ to use for adjusting function shape as in Figure 4. We use the shape of the function to tune system behavior and allow for problem-specific and scenario-specific adaptivity: B tunes how quickly R moves between 0 and 1, with lower values being more conservative (i.e., gradual), and ν adjusts where growth occurs on the curve. For completeness, note that limiting behaviors of R can capture the *no* mitigation and *pure* mitigation cases. For $\nu = 1$:

$$\lim_{B \to 0^+} R = 0 \tag{1}$$
$$\lim_{B \to \infty^+} R = 1 \tag{2}$$

In Limit 1, γ is always 0, which is equivalent to always applying full mitigation, i.e., always invoking a contingency. In Limit 2, γ is 1 until $t^c \leq 0$, which is equivalent to disabling mitigation, i.e., never activating a contingency until absolutely required. Under our mitigation strategy, even if the stopping region approximation is not conservative, it can be argued that the expected occurrence of disjointness violations is likely to be lower by virtue of the true stopping region volume growing less. Future work will formalize this argument.

IV. Demonstration: Automotive Active Safety Systems

Selective determinism decouples the navigation task into independent collision avoidance and guidance tasks. In an automobile, the collision avoidance task could be implemented as an active safety system and the guidance task assigned to a human driver. Because of vehicle inertial constraints, such a system would be at high risk for colliding reciprocal dances and thus is ideal for demonstrating our mitigation strategy.

We use the CARLA simulator [14] to implement a collision avoidance system on a human-guided vehicle. We conduct two demonstrations, a simple longitudinal trial that yields easy to interpret results, and a joint lateral/longitudinal trial to show behavior in more complex scenarios.

A. Longitudinal Active Safety

In these trials we have the human command the vehicle at full throttle into a stationary cyclist as shown in Figure 1. The intent with this scenario is to emulate a driver asleep at the wheel or distracted by a cell phone. For collision avoidance, we approximate vehicle motion with a constant acceleration model, which is simpler than the PhysX [15] model used by the simulator. We set the model deceleration to 90% of peak achievable simulation deceleration. This means the vehicle will typically slow more quickly than the collision avoidance system predicts. While this should result in conservative behavior, we nevertheless expect that the discrepancy in vehicle models will result in colliding reciprocal dance situations, even for this very simple scenario.

We compare three mitigation strategies:

1) **Constraint Tightening:** Nominal constraints are proximity-biased to contingency constraints.
2) **Conservative Deceleration:** The simplified motion model minimum is set to 80% of peak achievable.
3) **None:** The simplified motion model is used as-is.

Figure 6 shows speed profiles of the vehicle for each strategy, and Figure 5 shows computed throttle commands. The throttle oscillation in Figure 5 corresponds to the speed oscillation in Figure 6 toward the ends of the profiles.

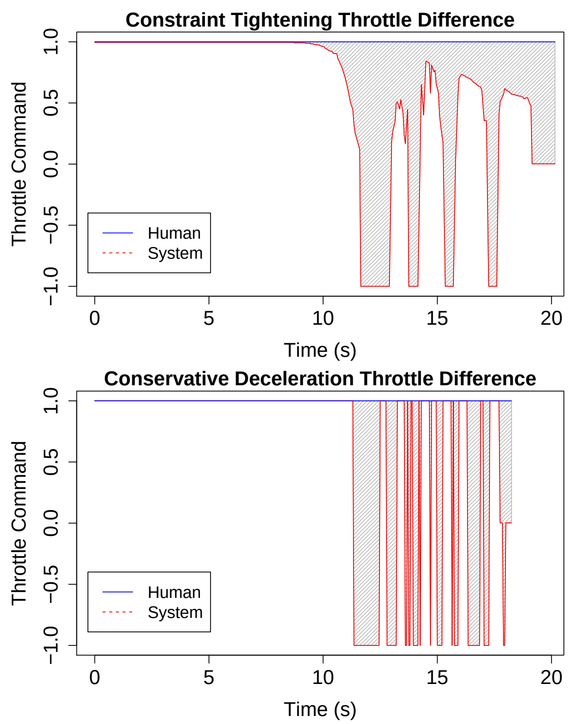

Fig. 5: Comparison of the guidance control (Human, blue) and collision avoidance controls (System, red) over time for Strategy 1 (top) and Strategy 2 (bottom). As a visual aid, the area between the curves is shaded. The reciprocal dance behavior is clear in the plots as the vehicle oscillates frequently between invoking and not invoking max deceleration contingency controls (and note that Strategy 1 significantly dampens the effect). All controls are in $[-1, 1]$. The red portion at the beginning of each indicates the vehicle idling and waiting for guidance input.

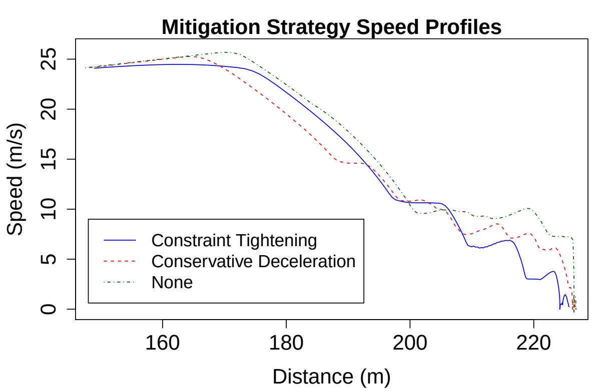

Fig. 6: Vehicle speed as a function of path position. The cyclist is located at approximately position $225m$. Only the portions of the profiles are shown where interaction effects begin to emerge; the prior portions are identical to each other.

These oscillations are due to the guidance signal disregarding interaction effects and the collision avoidance system repeatedly invoking contingencies. This is classic reciprocal dance behavior.

Strategy 1 permits the vehicle to maintain a higher speed along a greater extent of the path than Strategy 2 until approximately position $200m$, when it becomes more conservative. As the vehicle nears the cyclist, Strategy 1 keeps the vehicle slow and safely distant while Strategy 2 maintains relatively high speed until close proximity, only then to invoke significant deceleration. This demonstrates the utility of the proposed strategy: the adaptivity permits the vehicle greater dynamic range when it is safe, and more conservative dynamic range when needed. Finally, note that Strategy 3 exhibits a sharp speed decline at the end of the plot. This is caused by collision with the cyclist. In the absence of a mitigation strategy, the simple motion model resulted in a colliding reciprocal dance.

B. *Joint Lateral and Longitudinal Active Safety*

In these trials, we have the human command the vehicle at full throttle through moving traffic as shown in Figure 2. The intent with this scenario, like the first, is to emulate an inattentive or incapacitated driver. We adopt the same simplified longitudinal motion model in addition to a simplified steering model similar to that shown in Figure 3b. In this case it is not as straightforward to directly compare mitigation strategies due to the extra control dimension and the interactivity of the environment. Instead, we show how large a role the collision avoidance system plays in keeping the vehicle safe. Over multiple trials we observe that only those with the mitigation strategy active were able to consistently avoid collision.

Figure 7 shows how the collision avoidance system modifies the human guidance input in order to maintain safety. In the plots, the blue curve shows the human-provided guidance command that is taken as input to the system, and the red curve shows the control command that is output by the system to the vehicle. The large deviation of the red curves from the blue demonstrates that the collision avoidance system is doing a significant amount of work to keep the vehicle safe.

V. Conclusions & Future Work

The reciprocal dance problem is a common problem in mobile robotics systems that can also be dangerous in systems with inertial constraints. While exact solutions for avoiding the dance altogether are theoretically possible, we argue that no such solution is practical due to computational complexity. However, by formulating the problem in a principled way, we have derived a principled approach to mitigating it. The proposed mitigation strategy is beneficial over more naïve approaches because it provides adaptive behavior in order to maintain both safety and flexibility, and, from an implementation standpoint, can be relatively simple to use. The mitigation strategy was demonstrated in two scenarios, a longitudinal-only scenario and a joint lateral and longitudinal scenario.

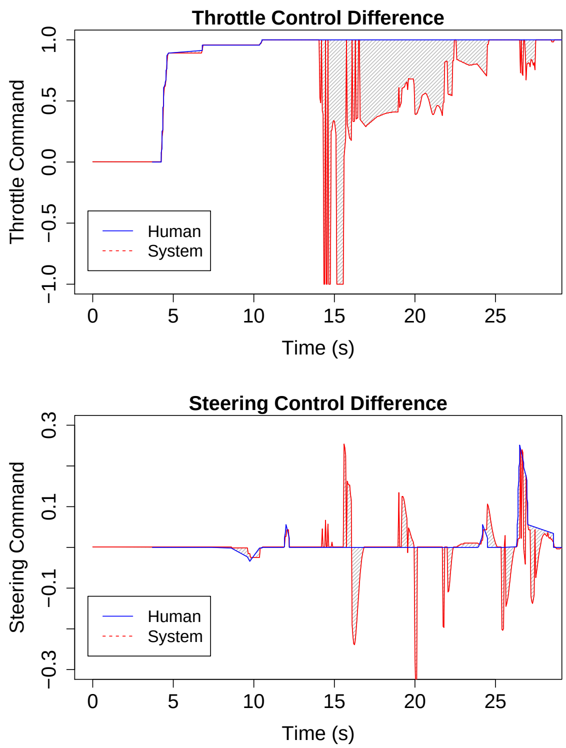

Fig. 7: Comparison of the guidance control (Human, blue) and collision avoidance controls (System, red) over time for the joint control demonstration. The area between the curves is shaded. All controls are in $[-1, 1]$. The red portions at the beginning indicate the vehicle waiting for guidance input.

Future work will further explore the verification and maintenance of Assumption 2 and additional argumentation that the mitigation strategy serves to reduce likelihood of collision. For the latter we believe the work in [16] can serve as a useful starting point. Also of critical importance is how to perform empirical verification of our system when interacting with intelligent agents in the physical world. This will allow us to quantify behavior on standard automotive test scenarios such as those described by NHTSA [17] and Euro NCAP [18]. We are also working to develop rigorous testing and engineering processes so that this system, and others like it, can be brought into conformance with industry safety standards such as UL 4600 [19] and ISO/PAS 21448 [20] that cover systems with autonomous capabilities.

References

[1] C. H. Papadimitriou and J. N. Tsitsiklis, "The Complexity of Markov Decision Processes," *Mathematics of Operations Research*, vol. 12, no. 3, pp. 441–450, 1987. [Online]. Available: http://www.jstor.org/stable/3689975

[2] D. S. Bernstein, R. Givan, N. Immerman, and S. Zilberstein, "The Complexity of Decentralized Control of Markov Decision Processes," *Math. Oper. Res.*, vol. 27, no. 4, pp. 819–840, 2002. [Online]. Available: http://dx.doi.org/10.1287/moor.27.4.819.297

[3] J. Goldsmith and M. Mundhenk, "Competition Adds Complexity," in *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6*, 2007, pp. 561–568. [Online]. Available: http://papers.nips.cc/paper/3163-competition-adds-complexity

[4] R. Emery-Montemerlo, G. J. Gordon, J. G. Schneider, and S. Thrun, "Approximate solutions for partially observable stochastic games with common payoffs," in *3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 19-23 August, New York, NY, USA*, 2004, pp. 136–143. [Online]. Available: http://doi.ieeecomputersociety.org/10.1109/AAMAS.2004.10017

[5] R. Pruim and I. Wegener, *Complexity Theory: Exploring the Limits of Efficient Algorithms*. Springer Berlin Heidelberg, 2005.

[6] R. Radner, "Team Decision Problems," *Ann. Math. Statist.*, vol. 33, no. 3, pp. 857–881, Sept 1962. [Online]. Available: http://dx.doi.org/10.1214/aoms/1177704455

[7] J. Marschak and R. Radner, *Economic Theory of Teams (Cowles Foundation Mon)*. Yale University Press, Feb 1972.

[8] S. Pendleton, T. Uthaicharoenpong, Z. J. Chong, J. G. M. Fu, B. Qin, W. Liu, X. Shen, Z. Weng, C. Kamin, M. A. Ang, L. T. Kuwae, K. A. Marczuk, H. Andersen, M. Feng, G. Butron, Z. Z. Chong, M. H. Ang, E. Frazzoli, and D. Rus, "Autonomous golf cars for public trial of mobility-on-demand service," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, Hamburg, Germany, September 28–October 2*, 2015, pp. 1164–1171. [Online]. Available: https://doi.org/10.1109/IROS.2015.7353517

[9] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robot. Automat. Mag.*, vol. 4, no. 1, pp. 23–33, 1997. [Online]. Available: https://doi.org/10.1109/100.580977

[10] P. Fiorini and Z. Shiller, "Motion Planning in Dynamic Environments Using Velocity Obstacles," *I. J. Robotic Res.*, vol. 17, no. 7, pp. 760–772, 1998. [Online]. Available: http://dx.doi.org/10.1177/027836499801700706

[11] J. P. van den Berg, J. Snape, S. J. Guy, and D. Manocha, "Reciprocal collision avoidance with acceleration-velocity obstacles," in *IEEE International Conference on Robotics and Automation, ICRA, Shanghai, China, 9–13 May*, 2011, pp. 3475–3482. [Online]. Available: http://dx.doi.org/10.1109/ICRA.2011.5980408

[12] J. K. Johnson, "Selective Determinism for Autonomous Navigation in Multi-agent Systems," Ph.D. dissertation, Indiana University, School of Informatics, Computing, and Engineering, September 2017. [Online]. Available: http://hdl.handle.net/2022/21670

[13] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios in a Frenét Frame," in *IEEE International Conference on Robotics and Automation, ICRA, Anchorage, Alaska, USA, 3–7 May*, 2010, pp. 987–993. [Online]. Available: https://doi.org/10.1109/ROBOT.2010.5509799

[14] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An Open Urban Driving Simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 1–16. [Online]. Available: http://carla.org

[15] NVIDIA Corporation, "PhysX," 2019. [Online]. Available: https://github.com/NVIDIAGameWorks/PhysX

[16] S. Karaman and E. Frazzoli, "High-speed flight in an ergodic forest," in *IEEE International Conference on Robotics and Automation, ICRA, 14-18 May, St. Paul, Minnesota, USA*, 2012, pp. 2899–2906. [Online]. Available: https://doi.org/10.1109/ICRA.2012.6225235

[17] W. G. Najm, J. D. Smith, and M. Yanagisawa, "Pre-Crash Scenario Typology for Crash Avoidance Research," U.S. Department of Transportation National Highway Traffic Safety Administration, April 2007, DOT-VNTSC-NHTSA-06-02. [Online]. Available: https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/pre-crash_scenario_typology-final_pdf_version_5-2-07.pdf

[18] Euro NCAP, "2018 Automated Driving Tests," 2018. [Online]. Available: https://www.euroncap.com/en/vehicle-safety/safety-campaigns/2018-automated-driving-tests/

[19] Underwriters Laboratories, "(forthcoming) Standard for Safety for the Evaluation of Autonomous Products," Underwriters Laboratories, Northbrook, IL USA, Standard, 2020. [Online]. Available: https://edge-case-research.com/ul4600/

[20] International Organization for Standardization, "Road vehicles – Safety of the intended functionality," International Organization for Standardization, Geneva, CH, Standard, January 2019. [Online]. Available: https://www.iso.org/standard/70939.html