Risk-Aware Decentralized Safe Control via Dynamic Responsibility Allocation (Student Abstract)

Yiwei Lyu[1*], **Wenhao Luo**[2], **John M. Dolan**[3]

[1]Department of Electrical and Computer Engineering, Carnegie Mellon University, USA
[2]Department of Computer Science, University of North Carolina at Charlotte, USA.
[3]Robotics Institute, Carnegie Mellon University, USA
yiweilyu@andrew.cmu.edu, wenhao.luo@uncc.edu, jdolan@andrew.cmu.edu

Abstract

In this work, we present a novel risk-aware decentralized Control Barrier Function (CBF)-based controller for multi-agent systems. The proposed decentralized controller is composed based on pairwise agent responsibility shares (a percentage), calculated from the risk evaluation of each individual agent faces in a multi-agent interaction environment. With our proposed CBF-inspired risk evaluation framework, the responsibility portions between pairwise agents are dynamically updated based on the relative risk they face. Our method allows agents with lower risk to enjoy a higher level of freedom in terms of a wider action space, and the agents exposed to higher risk are constrained more tightly on action spaces, and are therefore forced to proceed with caution.

Introduction

For multi-agent systems, centralized control could be computationally expensive and requires inter-agent communication. However, these conditions may not always be satisfied in the real world, which motivates the study on decentralized controller composition. Different approaches have been explored to translate centralized control into a decentralized setting, by splitting the constraints and separately solving individual optimization problems with split constraints, so that agents only need to make decisions based on local information, without the need to predict what others are going to do. (Wang, Ames, and Egerstedt 2016) partitions the constraints based on agents' various actuation limits. (Pierson et al. 2020) demonstrates how to divide the constraints given the known information of agent social personalities, egoistic vs. altruistic. In this work, we focus more on how to assign such agent identities based on risk evaluation. We consider risk caused by 1) possible collision among agents and 2) the uncertainty in agent motion, and propose to leverage the risk measurement information to compose decentralized safe controllers based on the responsibility share each agent is assigned, indicating the portion of constraint the agent is expected to respect compared to its pairwise companion. For each set of pairwise agents, a larger responsibility share is allocated to the agent facing less risk, and a smaller responsibility share is allocated to the agent with higher risk.

*Address: 5000 Forbes Avenue, Pittsburgh, PA 15213, USA.

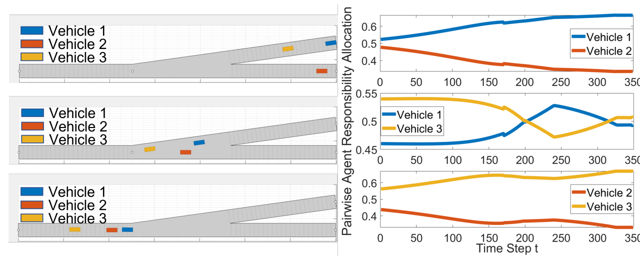

Figure 1: Dynamic responsibility allocation for pairwise agents in an autonomous driving ramp merging scenario, considering the risk from the possibility of collision and agent motion uncertainty in multi-agent interaction. The higher risk to which the agent is exposed, the less responsibility share is allocated to it, meaning the agent motion is more tightly constrained and forced to proceed with caution, compared to the agent facing lower risk.

The idea is to enforce tighter constraints on the motion of agent at higher risk, so that it has to proceed with more caution. Note that the goal of this work is to enable agents to make risk-aware decisions independently, instead of minimizing the entire system risk, which may make agents overconservative.

Control Barrier Function-inspired Risk Evaluation

Consider a multi-agent system with a total number of agents $N \in \mathcal{N}$. Since we are interested in pairwise agent safety, we define the pairwise safety function $h_{ij}(x)$ and safety set $\mathcal{H}_i$ as: $\mathcal{H}(x) = \{x \in \mathcal{X} : \ h_{ij}(x) = ||x_i - x_j||^2 - R_{safe}^2 \geq 0, \forall i \neq j\}$, $x_i, x_j \in \mathbb{R}^2$ for $i, j = \{1, ..., N\}$ are the positions of any pairwise agents i and j, and R_{safe} is the predefined safety margin.

Control Barrier Functions (CBF) (Ames et al. 2019) are used to define an admissible control space for safety assurance of dynamical systems. One of its important properties is its forward-invariance guarantee of a desired safety set. It is proved in (Ames et al. 2019) that for each pairwise agents, as long as they start inside the defined safe set and their control inputs belongs to the following admissible control space: $\mathcal{B}_{ij}(x) = \{u \in \mathcal{U} : \dot{h}_{ij}(x, u) \geq -\gamma(h_{ij}(x))$, then the system is always guaranteed to be collision-free. $\gamma \in \mathbb{R}^{\geq 0}$ is a CBF design parameter controlling system behaviors near the boundary of $h_{ij}(x) = 0$. Most existing works simply use CBF as a constraint for optimization-based controllers, serving as a binary verification of whether the system is still safe given the nominal control. In this work, we propose a CBF-inspired risk evaluation framework to characterize to what extent the system is safe or unsafe, in order to make the best use of the information CBF provides.

We define our pairwise safety loss function $L_{ij}(x, \bar{u})$ as: $L_{ij}(x, \bar{u}) = -\text{CVaR}_\alpha(\dot{h}_{ij}(x, \bar{u})) - \gamma h_{ij}(x) + c = -2(x_i - x_j)^T(\bar{u}_i - \bar{u}_j) - 2 \cdot \text{CVaR}_\alpha((x_i - x_j)^T(\epsilon_i - \epsilon_j)) - \gamma_i(||x_i - x_j||^2 - R^2_{safe}) + c$, where c as a constant offset is a large number to ensure $L_{ij}(x, \bar{u})$ is always positive to prevent unintended cancel-out when being accumulated later. $\bar{u}_i, \bar{u}_j \in \mathbb{R}^2$ are the agent's current velocities. $\epsilon_i, \epsilon_j \sim \mathcal{N}(\hat{\epsilon}, \Sigma)$ are random Gaussian variables with known mean $\hat{\epsilon} \in \mathbb{R}^2$ and variance $\Sigma \in \mathbb{R}^{2\times2}$, representing the uncertainty in each vehicle's motion. With the user-defined confidence level $\alpha \in (0, 1)$, Conditional Value at Risk calculated by $\text{CVaR}_\alpha(X) := \min_{z\in\mathbb{R}} \mathbb{E}\left[z + \frac{(X-z)^+}{1-\alpha}\right]$ (Rockafellar and Uryasev 2002), tells how bad the expected loss is if the condition is violated. We use $\text{CVaR}_\alpha(\cdot) \in \mathbb{R}$ to account for the worst-case scenario under motion uncertainty. It maps the risk of uncertainty to a real number and therefore is more suitable for embedding uncertainty information in risk measurement. The safety loss function $L_{ij}(x, \bar{u})$ represents given the user-specified confidence level α, how easily a safety violation could occur when agent i interacting with agent j. For a multi-agent system, the aggregated risk agent i faces posed by surrounding agents $R_i \in \mathbb{R}$ is therefore defined as: $R_i = \sum_{j=1}^{N} L_{ij}(x, \bar{u}), \quad \forall j \neq i$.

Decentralized Risk-aware CBF-based Controller

We know for any pairwise agents i and j, the centralized CBF-based safety constraint over agent velocity $u_i, u_j \in \mathcal{R}^2$ is in the linear form of: $\mathcal{B}(x_i, x_j) = \{u_i \in \mathcal{U}_i, u_j \in \mathcal{U}_j : A_i(u_i - u_j) \leq b\}$, where $A = -2(x_i - x_j)$ and $b = \gamma h(x_i, h_j) + 2 \cdot \text{CVaR}_\alpha((x_i - x_j)^T(\epsilon_i - \epsilon_j))$.

Theorem 1. *In a multi-agent system, agent safety during an interaction is formally guaranteed at a confidence level α, if for any pair of agents i and j, agent i takes the **Local Pairwise Responsibility Weight** $\omega_i = \frac{R_j}{R_i + R_j}$, so that the admissible control space in aforementioned centralized system is converted to: $\mathcal{B}(x) = \{u_i \in \mathcal{U}_i : A_i u_i \leq \omega_i b_i$ where $A_i = -2(x_i - x_j)^T \in \mathbb{R}^{1\times2}, b_i = \gamma h(x_i, x_j) + 2 \cdot \mathit{CVaR}_\alpha((x_i - x_j)^T(\epsilon_i - \epsilon_j)) \in \mathbb{R}$.*

The main idea behind the weight design of ω_i is to compare the relative level of risk each pairwise agents are exposed to, so that the agent with lower risk can enjoy a wider admissible control space by taking a larger responsibility portion, compared to agent with higher risk, which could be understood as it is already in a not that safe situation, and thus should only proceed with caution with a tighter safety bound. This design also reflects the idea of taking the neighbors of your neighbors into account, as for agent i's weight calculation, R_j embeds the constraints posed on agent j by its own neighbors. Proofs can be found in the Supplementary material. Therefore, for any pair of agents in multi-agent interaction, the **risk-aware CBF-based decentralized safe controller** is formulated as a quadratic program: $\min_{u_i \in \mathcal{U}_i} ||u_i - u'_i||^2, s.t \quad u_{min} \leq u_i \leq u_{max}, A_i u_i \leq \omega_i b_i$, where $u_i' \in \mathbb{R}^2$ is the nominal controller input, assumed to be computed by a higher-level task-related planner, u_{min} and u_{max} are the velocity bounds.

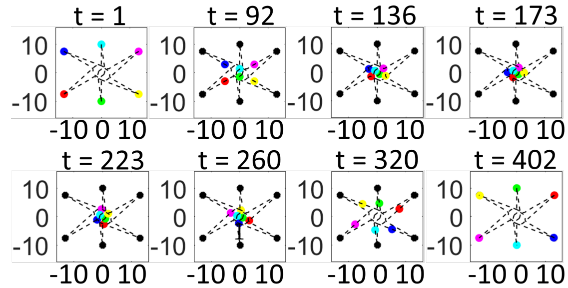

Figure 2: Multi-agent position swapping game with fixed and equal responsibility allocation (baseline method on left) v.s. dynamic responsibility allocation (proposed method on right).

Simulation

We demonstrate the validity and the effectiveness of the proposed method in an autonomous driving ramp merging scenario and a multi-agent position swapping game. In Fig. 1, we demonstrate that our proposed method is able to explicitly reason the relative risk level between pairwise vehicles in the presence of other traffic participants and adjust the responsibility shares accordingly to assign agents with lower risk a less tight constraint compared to agents with higher risk. In Fig. 2, compared to the baseline method where agent responsibility is simply equally assigned as a fixed value, our proposed method improves the overall task efficiency by 4.7% with less time spent.

References

Ames, A. D.; Coogan, S.; Egerstedt, M.; Notomista, G.; Sreenath, K.; and Tabuada, P. 2019. Control barrier functions: Theory and applications. In *18th European Control Conference (ECC)*.

Pierson, A.; Schwarting, W.; Karaman, S.; and Rus, D. 2020. Weighted buffered voronoi cells for distributed semi-cooperative behavior. In *IEEE international conference on robotics and automation (ICRA)*, 5611–5617.

Rockafellar, R. T.; and Uryasev, S. 2002. Conditional value-at-risk for general loss distributions. *Journal of banking & finance*, 26(7): 1443–1471.

Wang, L.; Ames, A.; and Egerstedt, M. 2016. Safety barrier certificates for heterogeneous multi-robot systems. In *American Control Conference (ACC)*, 5213–5218.