

Mutualistic Interactions in Heterogeneous Multi-Agent Systems

Alexander A. Nguyen[†], Faryar Jabbari[†], and Magnus Egerstedt[†]

Abstract—This paper presents a collaboration strategy that enables heterogeneous agents, i.e., different capabilities and dynamics, to accomplish tasks by working together. The collaboration between multiple agents is inspired by the ecological concept known as a mutualism, an interaction between two or more species that benefits everyone involved. A collaborative act is made possible through the composition of barrier functions, which allows the heterogeneous agents to work together safely. Moreover, a measure of collaborative potential is established to assess the merit of agents interacting with each other. Furthermore, the collaboration framework is provided for a general multi-agent setting. Finally, the collaboration framework’s efficacy is demonstrated in two case studies that necessitate collaboration between the heterogeneous agents to complete their respective tasks.

I. INTRODUCTION

Multi-agent systems have been utilized in a variety of applications, such as exploration [1], environmental monitoring [2], sensor coverage [3], and precision agriculture [4], which requires a team of agents to organize themselves efficiently and effectively to accomplish a shared objective [5].

This coordinated behavior can be achieved using homogeneous or heterogeneous teams. Typically, a homogeneous team will contain agents with similar body shapes and capabilities. Examples of homogeneous team cooperation include formation control [6], flocking [7], swarming [8], and foraging [9]. In contrast to this, a heterogeneous team will contain agents with dissimilar body shapes and capabilities, which can be leveraged to solve complex problems and increase task completion efficiency.

In this work, we will focus on heterogeneous agents. Heterogeneity can arise due to several factors, including body type [10], task allocation [11], hardware limitations [12], sensing capabilities [13], and safe operating regions [14] (also see the surveys in [15], [16]). It is advantageous for heterogeneous teams to exploit their diverse set of capabilities for collaboration purposes. In nature, for instance, we often see many examples of collaboration across species. In particular, symbiotic relationships can form between fundamentally different biological organism types. For example, a specific symbiotic relationship incentivizing collaboration is a mutualism [17], an interaction, e.g., exchange of nutrients or services, between two or more species where everyone benefits.

This work was sponsored in part by the US Office of Naval Research under grant number N00014-22-1-2625 and in part by the Graduate Assistance in Areas of National Need (GAANN) Fellowship under grant number P200A210021.

[†] A. A. Nguyen, F. Jabbari, and M. Egerstedt are with the Samueli School of Engineering, University of California, Irvine, Irvine, CA, 92697, USA. {alexaan2, fjabbari, magnus}@uci.edu.

Furthermore, this work will be considered in the context of engineered systems. Thus, it is important to ensure that safety requirements can be satisfied during the collaborative interactions of a heterogeneous multi-agent team. To this end, control barrier functions (CBFs) can be utilized as they provide a reliable means to guarantee the safety of control systems [18]. These CBFs must satisfy certain conditions such that a safe set can be rendered forward invariant. There are, for example, some approaches that can be used for constructing CBFs, which satisfy such conditions, including hand-design with careful consideration [19], [20] or synthesis using a learning-based framework [21], [22]. Moreover, selecting a suitable composition operator is crucial to obtaining a single, unified CBF in the presence of multiple CBF constraints. There have been previous approaches to define a composition operator when combining multiple CBFs in different problem settings, such as the use of Boolean logic formulas [23] and multiplication [24] operations.

The goal of this paper is to develop a collaboration framework for heterogeneous agents of an engineered system, which can exploit the functionalities possessed by every agent to accomplish tasks. In this work, the agents will collaborate to successfully reach a target location that is not reachable by an agent on its own. First, the heterogeneous agents will attempt to complete their tasks independently until the need for collaboration arises. Thus, a metric of collaborative potential, in the form of relative benefit or detriment, is defined to assess the merit of multiple agents undertaking a collaborative endeavor. Then, if collaboration is deemed beneficial, an agent in need will receive help from other agents nearby to accomplish their objective.

The remainder of this paper is organized as follows: Section II discusses how to encode pairwise interactions between agents within a shared workspace using barrier functions and establishes a suitable composition operator to combine multiple barrier functions for collaboration purposes. Section III formalizes a measure of the collaborative potential for agents to determine if collaboration would be beneficial or not. Section IV presents the collaboration framework for a general multi-agent setting. Section V provides numerical results in the form of two example scenarios to demonstrate the efficacy of the proposed collaborative control strategy. Section VI contains concluding remarks.

II. BARRIER FUNCTIONS FOR MULTI-AGENT SYSTEMS

This section discusses how pairwise interactions of multi-agent systems can be encoded using barrier functions. Furthermore, it describes a composition operator for combining multiple barrier functions.




Fig. 1: Agents with different mobility types, such as ground (portrayed as rabbits), amphibious (portrayed as turtles), and aerial (portrayed as birds), are contained in a shared workspace.

A. Encoding Pairwise Interactions Between Agents

Consider that N agents are to coexist in a shared workspace with states x_i belonging to d_i -dimensional manifolds \mathcal{X}_i , $\forall i \in \mathcal{N} = \{1, \dots, N\}$ (index set of the agents). Viewed in isolation, each agent is associated with a safe set, $\mathcal{S}_i \subseteq \mathcal{X}_i$, where they can operate safely and effectively in the absence of other agents.

As illustrated in Fig. 1, the safe set, \mathcal{S}_i , can be land, i.e., regions where ground-based agents (pictorially shown as rabbits) can operate or, conversely, the air, i.e., regions where aerial-based agents (pictorially shown as birds) can fly. \mathcal{S}_i can also be the union of several terrains, such as land and water regions, where amphibious agents (pictorially shown as turtles) can freely move in.

Assume that each such safe set, \mathcal{S}_i , is a superlevel set to a C^1 function, $h_i : \mathcal{X}_i \rightarrow \mathbb{R}$, such that $h_i(x_i) \geq 0 \Leftrightarrow x_i \in \mathcal{S}_i$ and, as such, h_i is defined to be a control barrier function (CBF) [19]. The CBF provides a safety certificate that ensures the chosen control input satisfies the following constraint

$$\frac{d}{dt}h_i(x_i) \geq -\alpha(h_i(x_i)), \quad (1)$$

for all times, given an extended class \mathcal{K}_∞ function, α , in order to render \mathcal{S}_i forward invariant [18].

The safe set, \mathcal{S}_i , can be interpreted as the safe operating region of agent i if there were no other agents present, which is similar to the notion of an ecological niche [25]. Evidently, the presence of other agents, such as robots, affects the ability to move in a shared workspace. For example, two agents typically cannot be at the same location at the same time to avoid collisions, which means the safe set would potentially get smaller for the agents.

However, it is also possible that the presence of other agents can expand the sets where the agents can operate safely. For example, a large robot could lift smaller robots onto ledges or carry them across obstructions. In other words, if we assume that these interactions are pairwise, then the change in the safe set for agent i at x_i , given the presence of agent j at x_j , can be encoded by the pairwise barrier function, $h_{ij}(x_i, x_j)$, which allows us to capture the influence of other agents' behaviors during pairwise

interactions within the workspace, as was done for obstacle avoidance scenarios in [26].

By combining the original barrier function, $h_i(x_i)$, with the new pairwise function, $h_{ij}(x_i, x_j)$, we get that agent i is safe, relative to agent j , if

$$H_i(x_i, x_j) = h_i(x_i) \oplus h_{ij}(x_i, x_j) \geq 0. \quad (2)$$

Here, \oplus denotes a composition operator that is left undefined for now. With this, safety for agent i is no longer just a function of x_i but also x_j .

As the focus of this paper is on how collaboration can enhance the performance of a team, we are interested in the best-case scenario, i.e., agent j is actively helping agent i . To that end, we can define the new safe set as

$$\mathcal{S}_{ij} = \{x_i \in \mathcal{X}_i \mid \exists x_j \in \mathcal{X}_j \text{ s.t. } H_i(x_i, x_j) \geq 0\}.$$

Here, the existential quantifier encodes the observation that safety is taken to mean that agent j can actively help agent i , represented by the inclusion of state x_j to define the new barrier function. Similar to before, we can draw a parallel to ecology by noticing that the shift from our original safe set, \mathcal{S}_i , to our new safe set, \mathcal{S}_{ij} , is similar to the concept of a niche expansion [27].

The pairwise scenario involving two agents can be extended to the multi-agent case in a straightforward manner, where the CBF of agent i is given by

$$H_i(x_1, \dots, x_N) = \bigoplus_{j=1}^N h_{ij}(x_i, x_j), \quad (3)$$

where we use the notational convention that $h_{ii}(x_i, x_i) = h_i(x_i)$.

Next, we will point out one possible way in which the composition operation can be formally defined in a multi-agent setting.

B. Composition of Barrier Functions

Let us now establish an appropriate definition for the composition operator \oplus . First, note that the composition of multiple CBFs has been investigated in various settings and for different types of applications.

In [23], Boolean composition, given as conjunctions (\wedge) and disjunctions (\vee), was shown to be equivalent to the min and max operators, respectively. However, these operators are not smooth, and the resulting composed CBF is no longer C^1 . This can be managed by using set-valued generalized gradients, but Boolean composition is not the appropriate route to pursue in the context of collaborative interactions.

If we, as before, let $H_i(x_i, x_j)$ be composed of the two CBFs $h_i(x_i)$ and $h_{ij}(x_i, x_j)$ and use conjunction, i.e., $H_i(x_i, x_j) = h_i(x_i) \wedge h_{ij}(x_i, x_j)$, then no expansion of the safe set is possible. In fact, by definition, $\mathcal{S}_{ij} \subset \mathcal{S}_i$ since $h_i(x_i) \geq 0$ has to hold for the composed CBF to be non-negative. Similarly, disjunction means that $\mathcal{S}_i \subset \mathcal{S}_{ij}$ which directly implies $H_i(x_i, x_j) = h_i(x_i) \vee h_{ij}(x_i, x_j) \geq 0$, i.e., $h_i(x_i) \geq 0$ or $h_{ij}(x_i, x_j) \geq 0$. As a result, non-safe interactions cannot be captured, which is not appropriate.

In [24], CBFs were, instead, composed through a multiplicative action, i.e., we let $H_i(x_i, x_j) = h_i(x_i) \oplus h_{ij}(x_i, x_j) = h_i(x_i) \cdot h_{ij}(x_i, x_j)$. The benefit of this construction is that composition is a smooth operator, as well as the intuitive fact that the “identity CBF” is equal to 1. In other words, if agent j has no impact on the safe set of agent i , then $h_{ij}(x_i, x_j) \equiv 1, \forall x_i, x_j$. This would work well if only two agents were present, but consider a situation with three agents where $h_i(x_i) \geq 0, h_{ij}(x_i, x_j) < 0, h_{ik}(x_i, x_k) < 0$. In that case, the composed CBF would satisfy

$$H_i(x_i, x_j, x_k) = h_i(x_i) \cdot h_{ij}(x_i, x_j) \cdot h_{ik}(x_i, x_k) \geq 0,$$

which is incorrect and problematic. The reason for this issue can be seen through collision avoidance. Assume that agent i will soon collide with agents j and k . Here, both of the pairwise CBFs are negative, but their composition is non-negative. Therefore, this doubly unsafe situation is being classified as safe.

In light of these potential obstructions, we chose to use an additive composition operator, meaning that

$$H_i(x_1, \dots, x_N) = \sum_{j=1}^N h_{ij}(x_i, x_j). \quad (4)$$

The resulting CBF is C^1 or piecewise C^1 , with all the constituent CBFs being C^1 or piecewise C^1 . Additionally, the identity CBF is given by $h_{ij}(x_i, x_j) \equiv 0, \forall x_i, x_j$, which is deemed safe since safety requires non-negativity rather than positivity.

However, a potential issue of using addition as a composition operator is that the signs of the constituent CBFs are no longer enough to determine the sign of the composite CBF since the magnitudes matter now as well. Although not problematic from a theoretical vantage point, this does mean that the constituent CBFs must be chosen carefully so that their magnitudes align properly to ensure what is classified as safe is indeed safe.

Lastly, note that the additive composition operator applies only to the constituent CBFs necessary for collaboration, $h_{ij}(x_i, x_j) \forall x_i, x_j$, resulting in the collaborative CBF constraint given by (4). Whereas the Boolean formulae composition operators, for example, unify multiple CBF constraints, such as the composed collaborative CBF and other non-collaborative CBFs, into a single constraint.

III. MEASURE OF COLLABORATIVE POTENTIAL

This section formalizes a measure of collaborative potential, given as relative benefit or detriment, in terms of safe sets to provide a way to characterize when agents would benefit from collaboration.

We defined, previously, a set-theoretic interpretation for how the safe set of an individual agent changes through the introduction of other agents, i.e., the new pairwise safe set, \mathcal{S}_{ij} , compared to the original safe set, \mathcal{S}_i . Now, we can ask if a collaborative endeavor is expected to be potentially

beneficial or not.¹ To this end, we will propose a measure for evaluating whether a pairwise interaction is expected to be beneficial or that a designer can use to determine how many heterogeneous agents should be deployed.

First, notice that if $\mathcal{S}_i \subset \mathcal{S}_{ij}$, then the presence of agent j is potentially beneficial to agent i , as long as the agents collaborate. However, it does not necessarily mean that agent j should collaborate solely because agent i benefits. Although, when $\mathcal{S}_i \subset \mathcal{S}_{ij}$, collaboration can help to achieve desirable results. For example, if agent i would like to reach a target, τ_i , where $\tau_i \notin \mathcal{S}_i$ while $\tau_i \in \mathcal{S}_{ij}$, then the second agent is necessary to accomplish the task at hand. Alternatively, when $\mathcal{S}_{ij} \subset \mathcal{S}_i$, the presence of the second agent is certainly detrimental to the first agent since, even in the best-case scenario, the safe set has shrunk.

Along these lines, a measure of size, $|\cdot|$, for the two sets would say something about the potential relative benefit or detriment during collaboration. If we define

$$\rho_{ij} = \frac{|\mathcal{S}_{ij}|}{|\mathcal{S}_i|}, \quad (5)$$

then $\rho_{ij} > 1$ indicates collaboration is potentially beneficial, whereas $\rho_{ij} < 1$ means it is potentially detrimental.

This discussion so far has been viewed through the lens of agent i . However, if collaboration is potentially beneficial for both agents, i.e., $\mathcal{S}_i \subset \mathcal{S}_{ij}$ and $\mathcal{S}_j \subset \mathcal{S}_{ji}$, then we have the possibility for what ecologists refer to as a mutualism. Namely, two agents of different “species” would team up and collaborate, thereby arriving at an arrangement benefiting both participants. If $\rho_{ij} > 1$ and $\rho_{ji} > 1$, the pairwise arrangement can be referred to as a “robot mutualism”.

Even in the absence of a mutualism, an altruistic act, where one agent is collaborating even though it receives a loss of benefit as a consequence, may still be useful. In nature, such acts are rarely observed across species since there is no evolutionary advantage associated with taking a loss of benefit for members of a different species. Instead, altruistic acts are observed within species as opposed to across species. In engineered systems, however, this distinction is not necessary.

There have been efforts made to understand when altruistic acts are expected to occur in nature. Hamilton’s Rule suggests that an altruistic act is worth undertaking if $rB > C$, where B is the benefit to the receiver, C is the cost to the provider, and r is the degree of kinship between the two organisms [28]. For instance, altruistic acts for offspring are typically more common and desirable from a survival vantage point, while the corresponding acts for strangers are less common.

Of course, in the design of an engineered system, one could encode the importance of agent i , or its tasks, to the overall mission through a scalar, κ_i , and define something analogous to Hamilton’s Rule for engineered agents, e.g., robots, that are collaborating.

¹For notational clarity, this discussion involves pairwise collaborations between two agents, but the developed concepts translate immediately to the N agent scenario.

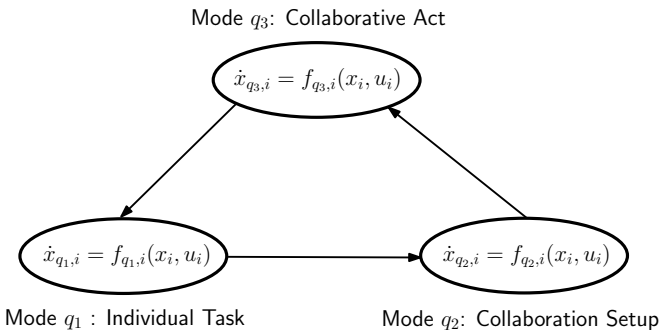


Fig. 2: Collaboration framework modes: individual task (q_1), collaboration setup (q_2), and collaborative act (q_3).

Such a rule could state that agent j should collaborate with agent i , even when it is detrimental to its performance, if

$$\kappa_i \rho_{ij} > \kappa_j \rho_{ji}, \quad (6)$$

where κ_i/κ_j is a scalar multiplier that encodes the importance of agent i 's performance improving relative to agent j 's, as observed in [29]. For example, a situation could arise where agent i has critical tasks to execute while agent j does not, so $\kappa_i/\kappa_j \gg 1$.

In the subsequent section, we will propose a method for enabling collaboration. Therefore, this paper answers two fundamental questions: "When is collaboration potentially beneficial?" and "How can such collaborations be achieved?". What we do not cover here, though, is the intermediary question of "Why should the agents collaborate?" which is left as future work.²

IV. COLLABORATION FRAMEWORK

This section outlines a framework for enabling heterogeneous agents to collaborate. First, we address how an agent can signal for help. Next, we define the collaboration framework's modes. Lastly, the system model and safety-critical controller are provided.

A. Event-Triggered Signal Condition

We consider heterogeneity in the context of safe operating regions, which means there may be areas in the workspace that an agent cannot reach independently.

For instance, agent i may need to traverse an unsafe region to complete their task, but it cannot due to an imminent safety violation. For such a scenario, an agent should be able to communicate to others nearby its need for help. In light of this, an event-triggered signal is developed so an agent can

²In ecology, a motivation for two or more species collaborating is the improvement of organisms' biological fitness, which corresponds to the ability of a species to survive and produce offspring successfully [30]. In robotics, the motivation for multiple agents to collaborate is not so clear. However, one could interpret their so-called "engineered fitness" as energy conservation, coverage quality, task completion, or maximization of artificial rewards. Ultimately, for a collaborative endeavor to be worthwhile, there must be a net positive benefit for the heterogeneous agents, regardless of whether the benefit is viewed as one-sided (altruism) or multi-sided (mutualism).

Mode q_1 : Individual Tasks

$$\begin{cases} \forall i \in \mathcal{N} : h_i(x_i) \geq 0, \\ \forall i \in \mathcal{N} : h_i(x_i) + \langle \nabla h_i(x_i), f_i(x_i, u_i) \rangle \Delta t \geq 0, \end{cases}$$

Mode q_2 : Collaboration Setup

$$\begin{cases} \forall i \in \mathcal{N} : h_i(x_i) \geq 0, \\ \exists i \in \mathcal{N} : h_i(x_i) + \langle \nabla h_i(x_i), f_i(x_i, u_i) \rangle \Delta t < 0, \end{cases}$$

Mode q_3 : Collaborative Act

$$\begin{cases} \forall i \in \mathcal{N} : H_i(x_1, \dots, x_N) \geq 0, \\ \exists i \in \mathcal{N} : h_i(x_i) + \min\{0, \langle \nabla h_i(x_i), f_i(x_i, u_i) \rangle \Delta t\} < 0. \end{cases}$$

Fig. 3: Transition conditions for each mode of the collaboration framework.

request help if it predicts danger, in the form of an unsafe region, is imminent over the time horizon Δt .

Let us define the evolution of agent i 's individual CBF, $h(x_i)$, over the time horizon Δt , as

$$h_i(x_i(t + \Delta t)) = h_i(x_i(t)) + \int_t^{t+\Delta t} \dot{h}_i(x_i(t')) dt', \quad (7)$$

where $\dot{h}_i(x_i(t')) = \langle \nabla h_i(x_i(t')), \dot{x}_i(t') \rangle$ is the individual CBF's time derivative, where agent i 's dynamics are defined in the general form $\dot{x}_i(t') = f_i(x_i(t'), u_i(t'))$ for now.

Assuming a short time horizon, we can use the forward Euler method for the numerical integration of ordinary differential equations. Therefore, with this in mind, the integral in (7) can be approximated as $\dot{h}_i(x_i(t))\Delta t$.

Thus, using the approximate evolution of agent i 's individual CBF, the event-triggered signal condition in continuous-time (CT) can be defined as $h_i(x_i(t + \Delta t)) < 0$ or as

$$h_i(x_i(t)) + \langle \nabla h_i(x_i(t)), f_i(x_i(t), u_i(t)) \rangle \Delta t < 0. \quad (8)$$

In practice, however, this CT signal is converted into a discrete-time (DT) signal during digital implementation. Therefore, after quantization, the event-triggered signal condition is redefined in DT as $h_i(x_i(kT + \Delta t)) < 0$ or as

$$h_i(x_i(k)) + \langle \nabla h_i(x_i(k)), f_i(x_i(k), u_i(k)) \rangle \Delta t < 0, \quad (9)$$

where the T in kT is dropped for ease of notation.

In this paper, we consider the time horizon to be one sampling period, i.e., $\Delta t = T$, but the extension to a larger time horizon is straightforward.

B. Collaboration Framework Modes

Heterogeneous multi-agent collaboration can be classified as a hybrid system since the agents' dynamics have the potential to change during collaborative interactions. As shown in Fig. 2, the collaboration framework depends on three modes to ensure heterogeneous agents can work together safely: individual task (q_1), collaboration setup (q_2), and collaborative act (q_3). Moreover, these modes will switch when the transition conditions, provided in Fig. 3, are satisfied at time t .

At first, each agent carries out its task independently until an agent predicts imminent danger using the event-triggered signal condition in (9). This prompts the agent in danger to request help, indicating a mode transition from individual task (q_1) into collaboration setup (q_2). Now, the agents initialize themselves for collaboration through a coordinated maneuver. Next, after the agents are configured properly, there is a mode transition from collaboration setup (q_2) into collaborative act (q_3). Here, the agent in need of help will receive assistance from other agents to accomplish its desired objective safely. Therefore, collaboration can only occur when there exists an agent safe through a pairwise interaction, but it would be unsafe on its own, i.e., $H_i(x_1, \dots, x_N) > 0$ but $h_i(x_i) < 0$. Then, once collaboration has concluded, the agents transition back into individual task mode to finish their objectives alone.

The dynamics of agent i were left in a general form for the event-triggered signal condition, but it should be defined as when an agent attempts to complete its task alone (mode q_1). Therefore, we set $f_i(x_i, u_i) = f_{q_1, i}(x_i, u_i)$ in Fig 3.

C. System Model

The dynamics of agent i are modeled in a control-affine form³, given by

$$\dot{x}_{q_v, i} = f_{q_v, i}(x_i) + g_{q_v, i}(x_i)u_i, \quad (10)$$

where $q_v \in \{q_1, q_2, q_3\}$ are the modes of the collaboration framework, $f_{q_v, i}$ and $g_{q_v, i}$ are locally Lipschitz, $x_i \in \mathcal{X}_i \subset \mathbb{R}^n$, and $u_i \in \mathcal{U} \subset \mathbb{R}^m$ is the set of admissible control inputs.

D. Quadratic Program

The collaboration framework is implemented using an optimization-based control scheme containing safety and non-safety constraints.

For agent i , we consider an actuator limitation constraint (non-safety) and a collaborative barrier certificate constraint (safety), given as

$$\frac{d}{dt}H_i(x) \geq -\alpha(H_i(x)), \quad (11)$$

where $\dot{H}_i(x) = L_f H_i(x) + L_g H_i(x)u$ is the time derivative of $H_i(x)$ defined using Lie derivative notation, i.e., $L_f H_i(x) = \nabla H_i(x) \cdot f_{q_v, i}(x_i)$ and $L_g H_i(x) = \nabla H_i(x) \cdot g_{q_v, i}(x_i)$, with $x = [x_1^\top, \dots, x_N^\top]^\top$.

Suppose we have a nominal controller that defines the agents' desired control strategy in a particular mode, q_v , of the collaboration framework without considering any safety constraints. As a result, each mode can have its behavior encoded through a possibly different nominal controller, $\hat{u}_{q_v, i}$, since an agent's dynamics can change due to the nature of collaboration.

³Many robotic systems exhibit control-affine dynamics when their models are derived using the Euler-Lagrange equations [31]




Fig. 4: Example scenarios for two agent collaboration.

Therefore, the nominal controller of agent i can take one of three modes, given as

$$\hat{u}_{q_v, i} = \begin{cases} \hat{u}_{q_1, i}, & \text{if mode } q_1 \text{ (individual task)} \\ \hat{u}_{q_2, i}, & \text{if mode } q_2 \text{ (collaboration setup)} \\ \hat{u}_{q_3, i}, & \text{if mode } q_3 \text{ (collaborative act)} \end{cases}$$

Now, a safety-critical controller can be used to guarantee the agents remain safe while attempting to track the nominal controller's reference signal in mode $q_v \in \{q_1, q_2, q_3\}$, given as the following Quadratic Program (QP)

$$\begin{aligned} u^* = \underset{u}{\operatorname{argmin}} \quad & \sum_{i=1}^N \|u_i - \hat{u}_{q_v, i}\|^2 & (12) \\ \text{s.t.} \quad & A_i(x)u \leq b_i(x), & \forall i \in \mathcal{N}, \\ & \|u_i\|_\infty \leq \bar{u}_i, & \forall i \in \mathcal{N}, \end{aligned}$$

where $A_i(x) = -L_g H_i(x)$, $b_i(x) = L_f H_i(x) + \alpha(H_i(x))$ are the linear constraints enabling collaborative interactions to occur between multiple agents and \bar{u}_i is an actuator limitation on agent i . Here, the decision variables, $u = [u_1^\top, \dots, u_N^\top]^\top$, are the agents' control inputs.

V. CASE STUDY: TWO AGENT COLLABORATION

This section showcases the collaboration framework in a two agent setting where the agents are motivated to work together by task completion⁴.

A. Background

The case studies for collaboration are conducted on example scenarios, illustrated in Fig. 4, which consider a team of $N = 2$ heterogeneous agents operating in a two-dimensional (2-D) domain.

By construction, a solution exists such that collaboration is feasible and the target state of agent i , τ_i , is reachable through a beneficial collaborative endeavor. That is, $\tau_i \notin \mathcal{S}_i$ while $\tau_i \in \mathcal{S}_{i,j}$. Furthermore, when utilizing the collaboration framework, each agent can complete its respective task successfully when working together. Recall that each mode, q_v , can have a different behavior, i.e., dynamics, encoded through the nominal controller's reference signal, $\hat{u}_{q_v, i}$, which switches based on the set of transition conditions being satisfied at time t .

For the case studies, we assume direct control authority over each agent's velocities through a single integrator dynamics model, which can be represented in control-affine

⁴The collaboration framework will work for the N agent scenario, but only two agents are needed to highlight its capability.

form as $f_{q_v,i}(x_i) = 0_{2 \times 1}$ and $g_{q_v,i}(x_i) = I_{2 \times 2}$ for all $q_v \in \{q_1, q_2, q_3\}$. The state vector of agent i is defined as $x_i = [p_{i,x}, p_{i,y}]^T$, where $p_{i,x}$ and $p_{i,y}$ are the x and y (planar) position states, respectively. The extended class \mathcal{K}_∞ function is chosen to be $\alpha(H_i(x)) = \gamma H_i(x)$, where $\gamma = 100$ and the sampling period is $T = 10^{-2}$ s. Lastly, the QP in (12) will be solved using the numerical convex optimization solver called CVXPY [32].

B. Example Scenario 1: Ground-Amphibious Robots

The first example, shown in Fig. 4(a), is the ground-amphibious robots scenario. The 2-D domain is partitioned into water (blue) and land (white) subdomains denoted as $\mathcal{D}_{\text{water}}$ and $\mathcal{D}_{\text{land}}$, respectively. Furthermore, the water subdomain has a width of ℓ_R , and its center point corresponds to the origin, $x_o = 0$.

The workspace is shared between a ground-based robot (rabbit) and an amphibious-based robot (turtle) that are modeled as point masses. The turtle can safely operate on land and water, i.e., $\mathcal{S}_t \subseteq \mathcal{D}_{\text{land}} \cup \mathcal{D}_{\text{water}}$, whereas the rabbit can safely operate on land only, i.e., $\mathcal{S}_r \subseteq \mathcal{D}_{\text{land}}$.

The individual and pairwise CBFs of the rabbit are constructed as

$$h_r(x_r) = p_{r,x}^2 - (\ell_R/2)^2, \quad (13)$$

$$h_{rt}(x_r, x_t) = \begin{cases} -h_r(x_t), & \text{if } x_r = x_t \\ 0, & \text{else} \end{cases}. \quad (14)$$

The individual CBF, $h_r(x_r)$, ensures the rabbit remains on land when carrying out its task, i.e., the rabbit is guaranteed to be more than half of the width away from the center of the water region in the x -direction. The pairwise CBF, $h_{rt}(x_r, x_t)$, allows the rabbit to traverse the water subdomain, with the turtle's assistance, during collaboration.

The individual and pairwise CBFs of the turtle are constructed as

$$h_t(x_t) = 0, \quad (15)$$

$$h_{tr}(x_r, x_t) = 0, \quad (16)$$

which are both defined as the identity CBF since the turtle can safely operate over the entire workspace and the rabbit has no pairwise influence on the turtle in this scenario.

Initially, the rabbit and turtle navigate toward their respective goal locations independently while in the individual task mode (q_1). Then, once the rabbit reaches the water subdomain boundary, it will signal the turtle for help. The robots now transition into the collaboration setup mode (q_2), where the rabbit and turtle will rendezvous near the intersection of the water and land subdomain boundaries. Then, after the robots coordinate themselves properly, the rabbit and turtle transition into the collaborative act mode (q_3), where the turtle will carry the rabbit across the water. Here, the rabbit benefits by getting ferried across the water, while the turtle benefits by enabling the rabbit to complete its task.

Fig. 5 shows the ground and amphibious robots always remain safe, i.e., the collaborative CBFs, $H_r(x_r, x_t)$ and




Fig. 5: Collaborative CBFs over time: (left) ground-amphibious robots and (right) square-triangle robots.

$H_t(x_r, x_t)$, are non-negative for all time. Furthermore, Fig. 6 highlights the different behaviors of the rabbit and turtle in each mode of the collaboration framework. In particular, Fig. 6(a), (b), and (c) correspond to the collaboration setup, collaborative act, and individual task modes, respectively.

C. Example Scenario 2: Square-Triangle Robots

The second example, shown in Fig. 4, is the square-triangle robots scenario. The 2-D domain consists of two horizontal surfaces of different heights connected by a line to form a step (black). Furthermore, the origin, $x_o = 0$, is defined at the bottom of the step.

The workspace is shared between two ground-based robots, denoted as square and triangle, that are modeled as convex polytopes defined as \mathcal{P}_s (square) and \mathcal{P}_t (triangle). Moreover, the triangle's and square's reference point, i.e., x_t and x_s , is located in the lower-left corner of their respective polygons. The square has a length of ℓ_s . The triangle has a base length of ℓ_t and a height of h .

Furthermore, the square and triangle have a single actuator generating velocity in a direction parallel to the bottom edge of each robot. This results in x -direction motion while both robots are in contact with the ground. However, when the square climbs on the triangle, it can generate displacement in the y -direction. Hence, we will set the triangle's height equal to the step's height, top horizontal surface, to ensure a smooth transition for the square when moving off the triangle.

Lastly, we initialize the square to the right of the triangle to ensure collaboration feasibility. For example, there would be no collaborative benefit if the square is placed to the left of the triangle. In this setting, the square's pairwise safe set is always smaller than the original safe set, i.e., $\mathcal{S}_{st} \subset \mathcal{S}_s$.

The individual and pairwise CBFs of the square are constructed as

$$h_s(x_s) = (p_{s,x} - \ell_t/2)^2 - (\ell_t/2)^2, \quad (17)$$

$$h_{st}(x_s, x_t) = \begin{cases} -h_s(x_s) + h - p_{s,y}, & \text{if } x_s \in \mathcal{P}_t \\ p_{t,x}(p_{t,x} - 2p_{s,x} + \ell_t), & \text{else} \end{cases}. \quad (18)$$

The individual CBF, $h_s(x_s)$, ensures the square avoids the triangle if it were stationary and placed at the origin, i.e., the square is guaranteed to not be within ℓ_t (triangle base length) to the right of the origin in the x -direction. The pairwise CBF, $h_{st}(x_s, x_t)$, has two purposes. First, it augments the individual CBF to include the triangle's position states so




Fig. 6: Simulation results for the ground-amphibious robots example scenario. The ground robot, i.e., rabbit (red), and amphibious robot, i.e., turtle (dark green), start at a location within their respective safe sets (turquoise stars). The blue and white regions represent water and land, respectively, while the green arrows represent the robots’ heading at time t . (a) The turtle moves toward the rabbit stuck at the boundary of the land subdomain to rendezvous. (b) The turtle carries the rabbit across the water subdomain to reach its target on the opposite land region. (c) The rabbit and turtle navigate toward their respective goal locations (purple stars) alone. [Supplemental Video: <https://youtu.be/bao5YV30dRk>].




Fig. 7: Simulation results for the square-triangle robots example scenario. The ground robots, i.e., square (dark green) and triangle (red), start at a location within their respective safe sets (turquoise stars). The black region represents the workspace boundaries, while the green arrows represent the robots’ heading at time t . (a) The square and triangle move toward the step to rendezvous. (b) The square climbs the triangle to reach the top of the step. (c) The square and triangle navigate toward their respective goal locations (purple stars) alone. [Supplemental Video: https://youtu.be/Aje6_oGUsA4].

the square avoids collisions with the triangle when both robots perform their tasks independently. Second, it allows the square to climb on the triangle without falling off to reach the step’s top level during collaboration.

The individual and pairwise CBFs of the triangle are constructed as

$$h_t(x_t) = p_{t,x}, \quad (19)$$

$$h_{ts}(x_s, x_t) = 0. \quad (20)$$

The individual CBF, $h_t(x_t)$, ensures the triangle remains on the lower horizontal surface, i.e., the triangle is guaranteed to stay right of the origin in the x -direction. The pairwise CBF, $h_{ts}(x_s, x_t)$, is the identity CBF since the square has no pairwise influence on the triangle.

Initially, the square and triangle navigate toward their respective goal locations independently while in the individual task mode (q_1). Then, once the square is unable to progress further, it will signal the triangle for help. The robots now transition into the collaboration setup mode (q_2), where the square and triangle will rendezvous at the step. Then, after the robots coordinate themselves properly, the square and triangle transition into the collaborative act mode (q_3), where

the triangle allows the square to reach the top horizontal surface. Here, the square benefits by using the triangle to climb onto the upper level of the step, while the triangle benefits by enabling the square to complete its task.

Fig. 5 shows the ground robots always remain safe, i.e., the collaborative CBFs, $H_s(x_s, x_t)$ and $H_t(x_s, x_t)$, are non-negative for all time. Furthermore, Fig. 7 highlights the different behaviors of the square and triangle in each mode of the collaboration framework. Fig. 7(a), (b), and (c) correspond to the collaboration setup, collaborative act, and individual task modes, respectively.

VI. CONCLUSIONS

In this paper, we developed a collaboration framework for a heterogeneous multi-agent team that has different capabilities and dynamics. The proposed approach enables an agent’s performance to improve by exploiting the functionalities of other agents in the workspace. However, before a collaborative interaction initiates, the agents should assess whether such an endeavor would be potentially beneficial or detrimental. In this work, we drew inspiration from ecology by considering the benefit to one “species” (altruism) or

multiple “species” (mutualism) during a collaborative interaction. Then, after the collaboration benefit is established, the heterogeneous agents can work together in a meaningful way to safely accomplish a shared objective using an optimization-based control scheme. Simulation results highlighted the collaboration framework’s capability through example scenarios in which two heterogeneous agents must work together to complete all tasks successfully.

VII. ACKNOWLEDGMENTS

The authors would like to thank Professor Jonathan N. Pauli and Mauriel Rodriguez Curras for helpful discussions about ecology.

REFERENCES

- [1] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, “Coordinated multi-robot exploration,” *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 376–386, May 2005.
- [2] M. Dunbabin and L. Marques, “Robots for environmental monitoring: Significant advancements and applications,” *IEEE Robotics and Automation Magazine*, vol. 19, no. 1, pp. 24–39, Feb. 2012.
- [3] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, “Coverage control for mobile sensing networks,” *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, Feb. 2004.
- [4] C. Ju and H. I. Son, “Multiple UAV systems for agricultural applications: Control, implementation, and evaluation,” *Electronics*, vol. 7, no. 9, Aug. 2018.
- [5] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.
- [6] P. Tabuada, G. J. Pappas, and P. Lima, “Feasible formations of multi-agent systems,” in *American Control Conference*, 2001, pp. 56–61.
- [7] K. Saulnier, D. Saldaña, A. Prorok, G. Pappas, and V. Kumar, “Resilient flocking for mobile robot teams,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1039–1049, Jan. 2017.
- [8] H. Shi, L. Wang, and T. Chu, “Swarming behavior of multi-agent systems,” *Journal of Control Theory and Applications*, vol. 2, no. 4, pp. 313–318, Nov. 2004.
- [9] D. A. Shell and M. J. Mataric, “On foraging strategies for large-scale multi-robot systems,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 2717–2723.
- [10] J. Roldan, P. Garcia-Aunon, M. Garzon, J. D. Leon, J. D. Cerro, and A. Barrientos, “Heterogeneous multi-robot system for mapping environmental variables of greenhouses,” *Sensors*, vol. 16, no. 7, July 2016.
- [11] S. Mayya, D. S. D’antonio, D. Saldaña, and V. Kumar, “Resilient task allocation in heterogeneous multi-robot systems,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1327–1334, Feb. 2021.
- [12] L. Yan, T. Stouraitis, and S. Vijayakumar, “Decentralized ability-aware adaptive control for multi-robot collaborative manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2311–2318, Feb. 2021.
- [13] J. Liu and R. K. Williams, “Monitoring over the long term: Intermittent deployment and sensing strategies for multi-robot teams,” in *IEEE International Conference on Robotics and Automation*, 2020, pp. 7733–7739.
- [14] S. Kim and M. Egerstedt, “Heterogeneous coverage control with mobility-based operating regions,” in *American Control Conference*, 2022, pp. 2148–2153.
- [15] Y. Rizk, M. Awad, and E. W. Tunstel, “Cooperative heterogeneous multi-robot systems: A survey,” *ACM Computing Surveys*, vol. 52, no. 2, pp. 1–31, Apr. 2020.
- [16] A. Prorok, M. Malencia, L. Carlone, G. S. Sukhatme, B. M. Sadler, and V. Kumar, “Beyond robustness: A taxonomy of approaches towards resilient multi-robot systems,” <https://arxiv.org/abs/2109.12343>, January 2021.
- [17] D. Boucher, S. James, and K. H. Keeler, “The ecology of mutualism,” *Annual Review of Ecology and Systematics*, vol. 13, no. 1, pp. 315–347, Nov. 1982.
- [18] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, “Control barrier functions: Theory and applications,” in *European Control Conference*, 2019, pp. 3420–3431.
- [19] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, Dec. 2016.
- [20] Q. Nguyen, A. Hereid, J. W. Grizzle, A. D. Ames, and K. Sreenath, “3D dynamic walking on stepping stones with control barrier functions,” in *IEEE Conference on Decision and Control*, 2016, pp. 827–834.
- [21] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, “Learning control barrier functions from expert demonstrations,” in *IEEE Conference on Decision and Control*, 2020, pp. 3717–3724.
- [22] M. Srinivasan, A. Dabholkar, S. Coogan, and P. A. Vela, “Synthesis of control barrier functions using a supervised machine learning approach,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 7139–7145.
- [23] P. Glotfelter, J. Cortés, and M. Egerstedt, “Nonsmooth barrier functions with applications to multi-robot systems,” *IEEE Control Systems Letters*, vol. 1, no. 2, pp. 310–315, June 2017.
- [24] L. Wang, A. D. Ames, and M. Egerstedt, “Multi-objective compositions for collision-free connectivity maintenance in teams of mobile robots,” in *IEEE Conference on Decision and Control*, 2016, pp. 2659–2664.
- [25] G. E. Hutchinson, “Concluding remarks,” *Cold Spring Harbor Symposia on Quantitative Biology*, vol. 22, pp. 415–427, 1957.
- [26] L. Wang, A. D. Ames, and M. Egerstedt, “Safety barrier certificates for collisions-free multirobot systems,” *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 661–674, Feb. 2017.
- [27] H. Sjödin, J. Ripa, and P. Lundberg, “Principles of niche expansion,” *Proceedings of the Royal Society B*, vol. 285, no. 1893, p. 20182603, Dec. 2018.
- [28] T. C. Bergstrom, “Evolution of social behavior: Individual and group selection,” *Journal of Economic Perspectives*, vol. 16, no. 2, pp. 67–88, June 2002.
- [29] M. Egerstedt, *Robot Ecology: Constraint-Based Design for Long-Duration Autonomy*. Princeton University Press, 2021.
- [30] S. O. Kimbrough, “The concepts of fitness and selection in evolutionary biology,” *Journal of Social and Biological Structures*, vol. 3, no. 2, pp. 149–170, Apr. 1980.
- [31] S. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [32] S. Diamond and S. Boyd, “CVXPY: A Python-embedded modeling language for convex optimization,” *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, Apr. 2016.