Leveraging Heterogeneous Capabilities in Multi-Agent Systems for Environmental Conflict Resolution

Michael E. Cao[1], Xinpei Ni[2], Jonas Warnke[2], Yunhai Han[2], Samuel Coogan[1], and Ye Zhao[2]

Abstract— In this paper, we introduce a high-level controller synthesis framework that enables teams of heterogeneous agents to assist each other in resolving environmental conflicts that appear at runtime. This conflict resolution method is built upon temporal-logic-based reactive synthesis to guarantee safety and task completion under specific environment assumptions. In heterogeneous multi-agent systems, every agent is expected to complete its own tasks in service of a global team objective. However, at runtime, an agent may encounter un-modeled obstacles (e.g., doors or walls) that prevent it from achieving its own task. To address this problem, we employ the capabilities of other heterogeneous agents to resolve the obstacle. A controller framework is proposed to redirect agents with the capability of resolving the appropriate obstacles to the required target when such a situation is detected. Three case studies involving a bipedal robot Digit and a quadcopter are used to evaluate the controller performance in action. Additionally, we implement the proposed framework on a physical multi-agent robotic system to demonstrate its viability for real world applications.

I. INTRODUCTION

Heterogeneous multi-agent systems with distinct mobility capabilities are generally capable of accommodating a larger variety of tasks than those consisting of a homogeneous team of agents [1]–[3]. To achieve autonomous team behaviors such as the multi-room patrolling shown in Figure 1b, a common approach is to automatically synthesize a controller for each agent, as this is often more efficient than crafting controllers by hand. However, creating controllers in this way has its own set of challenges, among which ensuring that the generated controllers do not cause any agents to perform tasks that would induce breakage or otherwise risk the agent's safety is a top priority [4].

In this paper, we study controller synthesis to resolve environmental conflicts such that multi-agent team specifications are fulfilled using the generalized reactivity (1) (GR(1)) fragment [6] of Linear Temporal Logic (LTL). The GR(1) fragment, in particular, allows for reactive synthesis algorithms that have favorable polynomial time complexity

while retaining the ability to encode a large variety of specifications [7]–[9]. These reactive synthesis methods are powerful because they provide formal guarantees on correct-by-construction controller behavior under any modeled action from the environment [10].

A major challenge of reactive synthesis methods is that they require an explicit model of the environment's capabilities and may not be robust to unexpected changes in these capabilities encountered at runtime [11]. For example, an unmodeled obstacle that interferes with the operation of the system may unexpectedly appear. Ensuring that the system's operation is robust to environmental changes is therefore an important area of research. To this end, multiple lines of research have been proposed in the literature, such as online synthesis of local strategies which are further patched to the original controller [12], offline analysis of counter-strategies to resolve unrealizable specifications [13], controller synthesis that can tolerate a finite sequence of environmental assumption violations [14], or robust metric automata designed such that the system state is maintained within a bounded ϵ-distance from the nominal state under unmodeled disturbances [15]. There also exist robustness methods that allow the system to identify specific broken environment assumptions [16]. However, none of these works studied the strategy of employing other agents to resolve environmental conflicts, which will be the focus of this paper, as this direction is underexplored. For example, the authors in [17] have studied the correction of broken assumptions, but focus on the cases where the broken assumption is due to unexpected behaviors by another agent operating within the workspace, which is resolved by changing that agent's behavior.

In heterogeneous teams, certain agents may have the capability to manipulate and correct a broken environment assumption for another agent. There have been works that explore the reconfiguration of the environment by the active agent in order to enable completion of its own objective [18]–[20]. these works, however, did not focus on employing the assistance of other agents. Motivated by this, our study focuses on leveraging each agent's individual capability to resolve broken environment assumptions that prevent another agent from achieving its objective. Formally, we consider scenarios where the broken environment assumption renders an agent's specification unrealizable, yet another agent has the ability to fix the violation.

The contribution of this paper is to (i) propose a navigation planning framework enabling agents in a multiagent team to leverage their heterogeneous capabilities to assist each other

[1]M. E. Cao and S. Coogan are with the School of Electrical and Computer Engineering, College of Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. {mcao34, sam.coogan}@gatech.edu

[2]J. Warnke, Y. Han, X. Ni and Y. Zhao are with the George W. Woodruff School of Mechanical Engineering, College of Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. {jwarnke3, xni32, yhan389, yzhao301}@gatech.edu
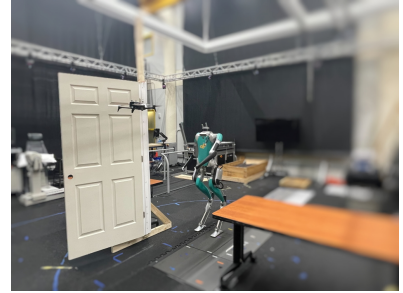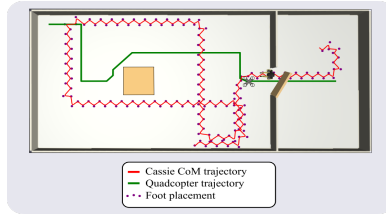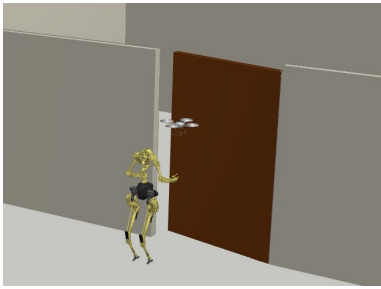
(a) Simulation set-up with a bipedal robot opening a door for the quadcopter

(b) A bird's eye view of Digit's and the quadcopter's trajectories

(c) Hardware set-up with a Digit walking robot and a quadcopter

Fig. 1: 3D rendering of case study 1 in Drake simulation [5]. The quadcopter approaches an unknown door obstacle preventing it from achieving its goal. The quadcopter then requests assistance from a bipedal walking robot Digit. Accordingly, Digit opens the door and both agents resume their original task. As the formulation is built on a correct-by-construction controller, the actions that Digit takes are guaranteed to be safe (barring detection of another unexpected obstacle), though they may not take the most direct path available. The controllers are then implemented in hardware experiments.

in resolving environmental obstacles using reactive synthesis; (ii) implement the proposed framework on a physical multi-agent robotic system composed of a Digit [21] bipedal robot and a Parrot Bebop 2 quadcopter. This is the first time that the LTL-based planning method is implemented on a heterogeneous robot team including a full-size humanoid robot for navigation tasks and obstacle resolution. The proposed framework has the following four main components:

- **Environment Characterization:** Observe the environment at runtime and verify whether the next state in the controller automaton would satisfy or violate any new specifications generated from runtime observations.
- **Safe Action Replanning:** Backtrack states in the automaton and replace an action that will lead to a safety violation with a known safe maneuver.
- **Violation Resolution:** Identify other agents with the capability of resolving the violation and assign one to be responsible for violation resolution.
- **Task Replanning:** Add the resolution of the violation to the assigned agent's objectives and trigger a change in behavior of that agent to execute its new objective.

We refer to these components collectively as the "coordination layer" that interacts with the other elements of the controller (see Figure 2).

The rest of the paper is outlined as follows: In Section II, we introduce the basics of LTL specifications and GR(1) formulas. We then formally define our problem statement in Section III. In Section IV, we provide an overview of the previous work that this study is built upon before detailing the main approach that enables heterogeneous cooperation in Section V. Section VI outlines several simulated case studies, and Section VII showcases our framework on a physical multi-agent system. Finally, we conclude the study and discuss potential future work in Section VIII.

II. PRELIMINARIES

In this study, we use the General Reactivity of Rank 1 (GR(1)) fragment [6] of Linear Temporal Logic (LTL) to specify desired tasks for each agent in a given environment.

GR(1) synthesis is used to automatically generate correct-by-construction finite state machines (FSM). The generated strategy is implemented as a two-player game between the agent and the environment, where the FSM guarantees the agent satisfies the goal and safety specifications for any modeled environment action [7], [22].

GR(1) allows for efficient synthesis while maintaining much of the expressiveness of LTL. In particular, GR(1) allows us to design temporal logic formulas (φ) with atomic propositions (AP) that can either be True ($\varphi \vee \neg\varphi$) or False ($\neg$True). With negation ($\neg$) and disjunction ($\vee$) one can also define the following operators: conjunction ($\wedge$), implication ($\Rightarrow$), and equivalence ($\Leftrightarrow$). There also exist temporal operators "next" ($\bigcirc$), "eventually" ($\lozenge$), and "always" ($\square$). Further details of GR(1) can be found in [6].

Our implementation uses the SLUGS reactive synthesis tool [23], which allows rules to be specified in a more human-intelligible structured slugs format using infix notation, non-negative integers, comparisons, and addition. These rules are automatically converted to GR(1) formulas which are used to synthesize a reactive controller.

III. PROBLEM FORMULATION

This paper studies a particular control synthesis problem where an agent in a heterogenous multi-agent team cannot complete its tasks because the environment has violated its assumptions at runtime. Formally, this occurs because an unmodeled environment behavior causes the specification to be unrealizable.

Let $\mathcal{P}$ denote the set of heterogeneous agents in a multi-agent team. When synthesizing a multi-agent controller, each agent $p \in \mathcal{P}$ within the team is given its own set of goal and safety specifications, denoted as φ_o^p and φ_s^p, respectively.

At the high-level, the environment is modeled using a coarse abstraction that divides the workspace into a set of N discrete regions $\mathcal{S} = \{s_0, s_1, ..., s_{N-1}\}$. As low-level controllers are responsible for planning agent actions within each coarse region, they can be swapped in and out to accommodate different agent types without largely affecting the high-level actions.

The set of known, irresolvable obstacles $\mathcal{O} \subset \mathcal{S}$ are accounted for as the set of safety specifications

$$\varphi_s^p := \bigwedge_{s \in \mathcal{O}} \square \neg s. \tag{1}$$

To account for the heterogeneity of the system, each agent $p \in \mathcal{P}$ is also modeled with an a priori known finite set of capabilities $C_p = \{c_{p0}, c_{p1}, c_{p2}, ...\}$. Examples of capabilities include "open doors", "inspect regions for hazards", or "climb stairs".

Obstacles that are resolvable but not known a priori are modeled as another subset $\mathcal{R} \subset \mathcal{S}$. Each resolvable obstacle $r \in \mathcal{R}$ has an associated action c_r and set of states S_r within which that action may be performed in order to resolve the obstacle and remove it from the environment. These properties are such that

$$c_r \in \bigcup_{p \in \mathcal{P}} C_p, \quad S_r \subset (\mathcal{S} \setminus \mathcal{O}). \tag{2}$$

Thus, an agent p is considered to be capable of resolving an obstacle r if $c_r \in C_p$. If an obstacle r whose corresponding resolution action is not present in any agent's capability set, the obstacle r will be considered as an unresolvable obstacle. By modelling agents and obstacles in this way, one can construct a robot team that is capable of resolving multiple types of obstacles by adding more robots with different capabilities to the team as needed.

It follows that any instance of an agent encountering an obstacle that it does not have the capability to resolve is considered a safety violation. We introduce an "augmented" set of safety specifications φ_a^p, which contains the same specifications as φ_s^p but in addition contains all of the additional safety specifications originating from unknown, resolvable obstacles:

$$\varphi_a^p := \varphi_s^p \bigwedge_{r \in R \ | \ c_r \notin C_p} \square \neg r \tag{3}$$

We can now formally define the problem statement.

Problem Statement: Assume a set of given controllers synthesized using φ_s^p, and a set of "actual" environment specifications φ_a^p such that one or more φ_o^p are unrealizable under φ_a^p. Once the system is detected to violate φ_a^p at runtime, create a generalizable formulation that assigns an agent p to resolve and remove conflicting specifications in φ_a^p such that the original synthesized controller satisfies φ_a^p.

IV. Controller Synthesis

To leverage the formal guarantees afforded by LTL, we synthesize navigation planners for each agent based on the planning framework detailed in [24]. In this section, we provide an overview of the task and motion planners, which serve as the foundation that the proposed coordination layer will be built on. In subsequent sections, we augment the high-level navigation planning structure to further encode collaborative behaviors capable of resolving environment assumption violations at runtime, which constitutes the main contribution.
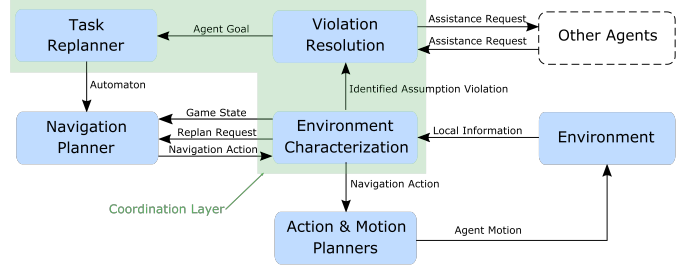

Fig. 2: Block diagram of collaborative task and motion planning framework. A coordination layer verifies that the desired actions generated by an offline-synthesized navigation planner are still safe based on environment information observed at runtime. If an environment assumption is violated, the navigation planner replans its current action to ensure the system enters a safe state while the coordination layer determines if the violation can be resolved by any other agent. The resolution is encoded and the plan progresses.

The approach from [7], [8], [22] adopted in [24] and used here is summarized as follows. To synthesize task planners, we construct two-player games between each agent and an abstracted environment. We automatically encode a variety of propositions about how the game may evolve within LTL specifications, including initialization assumptions, environment safety assumptions, system safety guarantees, and system liveness properties. We synthesize an automaton, represented as a finite state machine (FSM), that guarantees the agent will always win the game as long as all the environment assumptions hold true at runtime. During runtime, the current environment state is input to the FSM, which outputs an action for the agent. Each action is guaranteed to meet the safety specifications while bringing the agent closer to completing its task.

In this work we synthesize planners for a bipedal robot Digit [21] and a quadcopter to study heterogeneous autonomous multi-agent navigation. We accomplish this by constructing a two-player game between each agent and its environment, where we treat all other agents as part of a partially observable environment and encode the possible moves each agent can expect the others to make in environment safety specifications. Thus, we can easily expand to more than two agents while keeping the planner synthesization a two-player game. As the planner prioritizes safety, the optimality of the controller is not necessarily guaranteed. For full details, see [24].

V. Coordination Layer Design

In this section, we introduce the coordination layer on top of the synthesized controllers, containing the elements enabling agents to identify new safety specifications at runtime, replan actions when necessary, identify and assign other agents to resolve obstacles, and adjust the behavior of each agent to execute the conflict resolution. Throughout, we assume the agents possess the necessary capabilities to facilitate the described sensory and communication actions. An overview of the proposed planning framework is in

Figure 2.

A. Environment Characterization

At runtime, the Environment Characterization (EC) block passes the game state abstracted from the environment to the Navigation Planner (NP) block and requests a navigation action. The EC block determines new safety specifications based on the observed environment and verifies if the navigation action would violate these specifications. This block is the main component responsible for observing the specifications in φ_a^p that were not known during synthesis. If a safety violation occurs, the EC block signals the NP block to replan its current action.

B. Safe Action Replanning

The safe action replanning occurs within the NP block. At each step, the current environment state is fed to the FSM to generate a correct system action. When the replanning flag is raised, the navigation planner backtracks to the previous state in the FSM and extracts a new system action to avoid the safety violation. The new action is passed to the EC block, which passes it on to the lower-level planners if deemed safe.

C. Violation Resolution

When the EC block determines that the safety specifications based on the observed environment do not match the safety specifications used during offline synthesis, it also passes the details of the violation to the Violation Resolution (VR) block, including the action(s) c_r required, as well as the state(s) S_r at which those actions must be performed in to resolve the obstacle. The VR block then identifies which agent $p \in \mathcal{P}$ has the capability to resolve the obstacle, i.e., any agent p such that $c_r \in C_p$. This component then broadcasts a request for an appropriate agent to assist.

D. Task Replanning

The Task Replanner (TR) block receives incoming requests for assistance in the form of updated system goals and is responsible for adjusting the agent's strategy to assist the agent in need. Once a violation has been detected and assisting agents have been assigned, the TR block can trigger a resynthesis of each of the affected agent's controllers with their new objectives. The helping agent is assigned to resolving the obstacle, and the requesting agent is assigned to a known safety objective that does not interfere with the helping agent. After the controller detects that the obstacle has been resolved, another resynthesis is triggered, returning the agents to their original objectives.

This method stores previous objectives in a stack; if a new resolvable obstacle is encountered while resolving a known obstacle, the resynthesis targets the new obstacle and pushes the previous set of objectives onto the stack. Once an obstacle is resolved, the top set of objectives in the stack is popped off, and controllers are resynthesized to these objectives. Beyond changing the target locations in the original objective, any number of new goal tasks or locations can be encoded in the new specifications to assist

another agent that has encountered an assumption violation. Additionally, the agent that has encountered an obstacle that it cannot resolve on its own can be assigned a new task to complete while it waits for the obstacle to be resolved. A detailed pipeline of this process is shown in Algorithm 1.

It should be noted that the algorithm is built upon an assumption of the existence of a sequence of obstacle resolutions that eventually results in the original objectives becoming achievable. In cases where this is not true, i.e., no agent exists that can resolve the obstacle or a set of resolution objectives are pushed onto the stack recursively, it would be impossible for the agents to accomplish their original set of objectives without external intervention. In such cases, synthesis or resynthesis of the automaton will fail, signaling the need for human operator intervention. One may avoid these cases by introducing agents with new capabilities.

Algorithm 1: Resynthesis for Conflict Resolution

---

for $p \in \mathcal{P}$ **do**
  | $\varphi_o^p \leftarrow$ initial objectives;
end
$\varphi_{\text{stack}} \leftarrow \emptyset$;
synthesize controllers;
while system active **do**
  | execute controllers;
  | **if** resolvable obstacle r encountered by agent p_r at s_o **then**
  | | **if** $c_r \notin C_{pr}$ **then**
  | | | push all of current φ_o^p onto φ_{stack};
  | | | change φ_o^{Pr} to safe behavior;
  | | | **for** $p \in \mathcal{P}$ **do**
  | | | | **if** $c_r \in C_p$ **then**
  | | | | | add state in S_r to φ_o^p;
  | | | | | synthesize controllers;
  | | | | | break;
  | | | **end**
  | | | continue;
  | | **else if** $r \in \varphi_o^{Pr}$ **then**
  | | | resolve obstacle r;
  | | | pop φ_o^p off of φ_{stack};
  | | | synthesize controllers;
  | | | continue;
  | | **end**
  | **end**
end

Notice that since the controllers are synthesized by GR(1) reactive synthesis, they are correct by construction [7].

We additionally develop an alternative method involving an additional runtime-assignable goal location specified in the assisting agent's liveness specifications, which enables a limited amount of coordination without resynthesis. However, details of this method are omitted due to space constraints. Instead, for full implementations of either method, we point the reader to the GitHub repository located at `https://github.com/GTLIDAR/safe-nav-locomotion`.

VI. CASE STUDIES

In this section, we implement and evaluate the framework described in Section IV, which was primarily built to synthesize controllers for the bipedal walking robot platform Digit, designed by Agility Robotics [21]. We consider a second quadcopter agent in the environment that has dramatically

different capabilities in both mobility and manipulation. A quadcopter, which lacks the ability to manipulate objects, is instead able to perform maneuvers that are unavailable to Digit, such as backward movement or $180°$ turns in a single region. In this study, the quadcopter is assumed to fly above Digit at all times, so that collision is not a concern. This assumption is made to reduce the complexity of the controller synthesis for these case studies, as employing collision avoidance, while trivial from an implementation standpoint thanks to [24], drastically increases the computation time. Additionally, we preserve the belief space planning framework proposed in [24], [25], allowing Digit to infer the quadcopter location if it is out of Digit's visible range. Digit's locomotion planner is designed based on the phase-space planning framework in [26].

Resolvable obstacles are also implemented within the simulated environment. We recall from Section III that each resolvable obstacle r has an associated state S_r where an action c_r must be performed by an agent p for which $c_r \in C_p$ in order to resolve the obstacle and remove it from the environment. These properties are directly inserted into the simulation environment. The set of resolvable obstacles $\mathcal{R}$ and the set of agents and their capabilities C_p are implemented as separate dictionary data structures. As such, this framework is generalizable as one would simply need to add the appropriate agent capabilities and obstacle resolutions to each dictionary.

To represent the potential for obstacles to appear at runtime, two separate simulated environments are initialized, one without the resolvable obstacle. This instance of the environment, representing φ_s^p, is used for the initial synthesis, and then the resulting controller is applied to the environment instance containing the resolvable obstacles, which corresponds to φ_a^p. At runtime, if an agent enters a state containing a resolvable obstacle r that it is unable to resolve (i.e., violates φ_a^p), the controller is able to check that a violation has happened in the simulated environment, and sends this information to the simulation to assign new objectives to each agent accordingly, such that the agent p tasked with resolving the obstacle fulfills $c_r \in C_p$. Thus, the simulation running each of the controllers is responsible for the VR and TR blocks of the coordination layer, while the separate simulated environment instances simulate the EC component. The Safe Action Replanner is built directly into the NP block.

Three case studies utilizing the synthesized controllers are presented to evaluate the proposed approach. For each case study, an environment is created where a quadcopter and Digit are each running on their own controller and have their own task objectives to complete. The environment is abstracted into a 7×13 coarse set of regions $\mathcal{S} = \{s_0, s_1, ..., s_{90}\}$ such that s_0 is the northwestern-most region and increments following English reading orientation (i.e. incrementing left to right, then starting at the leftmost region on the next row). This setup can be seen in Figures 3-4.

We consider a team of agents $P = \{\texttt{quadcopter and Digit}\}$ with unique capabilities $C_{\text{quad}} := \{\texttt{sense}\}$, $C_{\text{Digit}} :=$

$\{\texttt{push}\}$. Resolvable obstacles that may appear within a region s_i in the environment consist of two types: $r = \texttt{door}$ and $r = \texttt{uncertainty}$. A resolvable obstacle of type $r = \texttt{door}$, if found in s_i, has properties

$$S_r = \{s_i\}, \ c_r = \texttt{push}, \quad (4)$$

and represents physical doors that the quadcopter cannot fly through, but are able to be opened by Digit. Resolvable obstacles of type $r = \texttt{uncertainty}$ have the properties

$$S_r = \{s_{\text{north}}, s_{\text{east}}, s_{\text{south}}, s_{\text{west}}\}, \ c_r = \texttt{sense}, \quad (5)$$

and represent regions in which Digit is uncertain about its capabilities to safely traverse through the environment, but the quadcopter is able to scout them by visiting any adjacent region.

We design a set of objective specifications for each agent $p \in \mathcal{P}$ such that the agent alternates between visiting two regions in the environment $s_A, s_B \in \mathcal{S}$, with an AP $\texttt{scout}$ unique to each agent, which initializes to False, to track which region the agent should head towards. The set of objective specifications is thus

$$
\begin{aligned}
\text{Patrol}_p(s_A, s_B) = \quad & (6) \\
\Box \Diamond (s_A \wedge \neg \texttt{scout}_p) & \\
\wedge \Box((s_A \wedge \neg \texttt{scout}_p) \Rightarrow \bigcirc(\texttt{scout}_p \wedge \neg s_A)) & \\
\wedge \Box \Diamond (s_B \wedge \texttt{scout}_p) & \\
\wedge \Box((s_B \wedge \texttt{scout}_p) \Rightarrow \bigcirc(\neg \texttt{scout}_p \wedge \neg s_B)) & \\
\wedge \Box \neg (s_A \wedge \texttt{scout}_p) & \\
\wedge \Box \neg (s_B \wedge \neg \texttt{scout}_p). &
\end{aligned}
$$

While the figures in this section mainly feature abstractions of the environment in order to easily illustrate the behaviors of each agent, the computed control actions are applicable to a real 3D simulation environment, as shown in Figure 1. Those two subfigures show the real-world interactions resulting from the behavior in the first case study.

### A. Case Study 1: Opening A Door

The first case study leverages Digit's manipulation capability in the environment with higher dexterity and power than the quadcopter. The quadcopter is tasked with patrolling between s_{homeQ} and s_{awayQ}, where s_{homeQ} is a region in the left room and s_{awayQ} is in the right room. Digit is tasked with patrolling between s_{homeC} and s_{awayC}, where both s_{homeC} and s_{awayC} are in the left room:

$$
\begin{aligned}
\varphi_o^{\text{quad}} &:= \text{Patrol}_{\text{quad}}(s_{\text{homeQ}}, s_{\text{awayQ}}), \quad (7) \\
\varphi_o^{\text{Digit}} &:= \text{Patrol}_{\text{Digit}}(s_{\text{homeC}}, s_{\text{awayC}}).
\end{aligned}
$$

At runtime, the quadcopter discovers an obstruction at $s_{\text{door}} = s_{47}$ while in $s_{\text{safe}} = s_{46}$ that prevents it from accomplishing its objective in the form of a closed door that cannot be flown through but can be opened by Digit. A resynthesis of objectives is triggered, where the quadcopter is now tasked with hovering outside the door, and Digit is
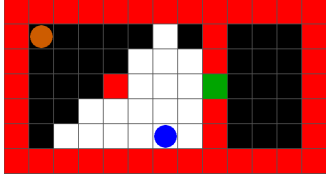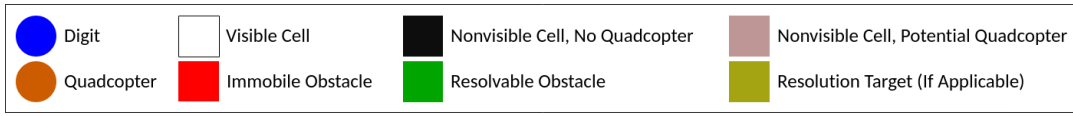
(a) Initial Configuration
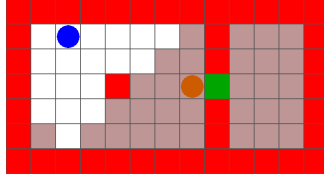
(b) The quadcopter's controller encounters a door, so instead it elects to have the quadcopter wait in the preceding region.
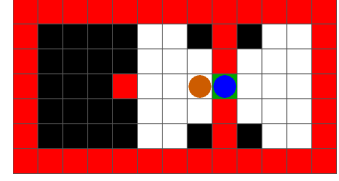
(c) Digit opens the door and the original objectives are reinstated. Both agents are now able to complete their objectives.

Fig. 3: Execution of case study 1 leveraging Digit's manipulation abilities to clear the path for the quadcopter. Digit's objective is patrolling the left room, while the quadcopter's objective is delivering to the right room. However, the quadcopter discovers a closed door separating the two rooms at runtime, prompting Digit to navigate to open it so that both agents are able to complete their objectives.

tasked with visiting one of its initial patrol points and the closed door, resulting in

$$\varphi_o^{\text{quad}} := \text{Patrol}_{\text{quad}}(s_{\text{safe}}, s_{\text{safe}}), \quad (8)$$
$$\varphi_o^{\text{Digit}} := \text{Patrol}_{\text{Digit}}(s_{\text{homeC}}, s_{\text{door}}).$$

Once Digit visits the door, it is considered open and the obstacle is removed from the environment, triggering another resynthesis which returns both agents to their original target objectives. A walkthrough of the execution of this case study is shown in Figure 3. For this case study, we also used low-level planners to generate safe motions for the quadcopter and Digit, including center of mass (CoM) trajectories and foot placements. A 3D visualization can be seen in Figure 1 and the included video.

B. Case Study 2: Scouting Ahead

The second case study involves Digit encountering several states and not knowing whether each state is safe to traverse on foot, requiring the help of the quadcopter's heightened sensing capabilities. To this end, the quadcopter is set to patrol between s_{homeQ} and s_{awayQ}, where s_{homeQ} and s_{awayQ} are in the left room, while Digit must patrol between s_{homeC} and s_{awayC}, where s_{homeC} is a region in the left room and s_{awayC} is in the right room:

$$\varphi_o^{\text{quad}} := \text{Patrol}_{\text{quad}}(s_{\text{homeQ}}, s_{\text{awayQ}}), \quad (9)$$
$$\varphi_o^{\text{Digit}} := \text{Patrol}_{\text{Digit}}(s_{\text{homeC}}, s_{\text{awayC}}).$$

At runtime, Digit encounters region $s_{\text{uncertain1}} = s_{34}$ while at $s_{\text{safe1}} = s_{33}$, and it is unsure about its ability to traverse this region. A resynthesis is triggered, where the quadcopter is tasked with observing the uncertain region by visiting any of the adjacent regions (in this case, we select the region $s_{\text{uncertain1W}}$ directly west of $s_{\text{uncertain1}}$), resulting in

$$\varphi_o^{\text{quad}} := \text{Patrol}_{\text{quad}}(s_{\text{homeQ}}, s_{\text{uncertain1W}}), \quad (10)$$
$$\varphi_o^{\text{Digit}} := \text{Patrol}_{\text{Digit}}(s_{\text{safe1}}, s_{\text{safe1}}).$$

Once the quadcopter observes the unknown region, the resolvable obstacle is removed from the environment. If the

quadcopter senses that the region is not traversable by Digit, then the region is added to O and will be considered as an immovable obstacle during future synthesis. In this specific case study, $s_{\text{uncertain1W}}$ is found to be untraversable.

The two agents are returned to their original objectives, outlined in (9), before Digit encounters another uncertain state at $s_{\text{uncertain2}} = s_{60}$ while in $s_{\text{safe2}} = s_{59}$, where the process repeats. The quadcopter is tasked with visiting $s_{\text{uncertain2W}}$, directly west of $s_{\text{uncertain2}}$, while Digit is instructed to stay at s_{safe2}. The uncertain region is found to be traversable by Digit, and both agents are returned to their original objectives, now able to fulfill them. A walkthrough of the execution of this case study is shown in Figure 4.

C. Case Study 3: Chain of Conflicts

The third case study merges the previous two and requires both agents to resolve an obstacle. For this case study, the quadcopter encounters a door while on its way to resolving an uncertain region encountered by Digit, requiring Digit to first open the door, thus demonstrating the coordination layer's ability to handle multiple resolvable obstacles in a chain when required.

Initially, Digit is tasked with patrolling between s_{homeC} in the leftmost room and s_{awayC} in the center room, while the quadcopter is tasked with patrolling between s_{homeQ} and s_{awayQ}, both in the rightmost room:

$$\varphi_o^{\text{quad}} := \text{Patrol}_{\text{quad}}(s_{\text{homeQ}}, s_{\text{awayQ}}), \quad (11)$$
$$\varphi_o^{\text{Digit}} := \text{Patrol}_{\text{Digit}}(s_{\text{homeC}}, s_{\text{awayC}})$$

Digit encounters an uncertain state at $s_{\text{uncertain}} = s_{43}$ while in $s_{\text{safeC}} = s_{43}$, triggering a resynthesis requiring the quadcopter to sense the true traversibility of that state by visiting $s_{\text{uncertainE}} = s_{44}$, resulting in

$$\varphi_o^{\text{quad}} := \text{Patrol}_{\text{quad}}(s_{\text{homeQ}}, s_{\text{uncertainE}}), \quad (12)$$
$$\varphi_o^{\text{Digit}} := \text{Patrol}_{\text{Digit}}(s_{\text{safeC}}, s_{\text{safeC}}).$$

However, the quadcopter encounters a closed door at $s_{\text{door}} = s_{47}$ while in $s_{\text{safeQ}} = s_{48}$ on its way to resolve Digit's uncertainty, triggering another resynthesis where
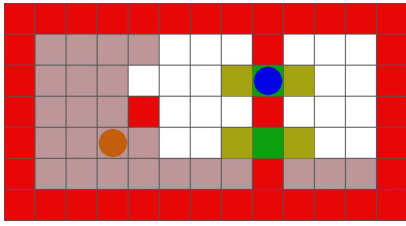
(a) Whenever Digit's controller encounters an uncertain region, it has Digit wait in the preceding region.
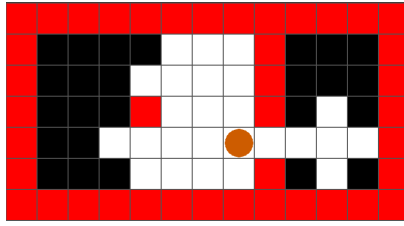
(b) The quadcopter senses that the first region is an obstacle, but the second region is traversable by Digit.
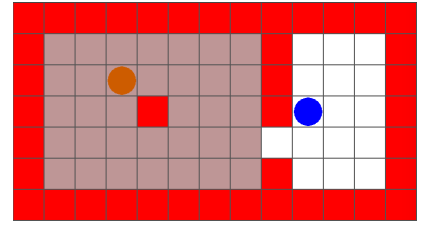
(c) Original objectives are reinstated, with both agents now able to meet their objectives.

Fig. 4: Execution of case study 2 leveraging the quadcopter's sensory capabilities to find a traversable path for Digit. The quadcopter's objective is to patrol the left room, while Digit's objective is to deliver something to the right room. However, Digit encounters several uncertain regions, requiring the quadcopter to observe each until it finds a traversable region.

Digit is tasked with opening the door, resulting in $\varphi_o^{\text{quad}} := \text{Patrol}_{\text{quad}}(s_{\text{safeQ}}, s_{\text{safeQ}})$, $\varphi_o^{\text{Digit}} := \text{Patrol}_{\text{Digit}}(s_{\text{safeC}}, s_{\text{door}})$. Once the door is opened and resolved, the quadcopter travels to the uncertain state and resolves that obstacle, resulting in both agents being able to accomplish their objectives.

For this case study, in lieu of a figure illustrating the abstracted simulation, we apply the computed control actions to a physical robotic system consisting of a bipedal robot and a quadcopter, detailed in the next section.

VII. HARDWARE IMPLEMENTATION

We replicated case study 3 (see Sec. VI-C) on physical hardware, which involved coordination between a bipedal robot Digit and a quadcopter Parrot Bebop 2. A video of the experiment is included as supplementary material.

Digit has an upper body that consists of a torso and two arms with 4 actuatable joints on each, which allows it to achieve manipulation skills and enables it to open a door during the hardware experiment (see Fig. 5). For this experiment, we use the Digit controller provided by Agility Robotics, which takes the 2D world waypoint position as input and automatically plans Digit's walking gait to approach the desired position. Hence, at runtime, Digit is commanded to track the waypoints sent by the task and motion planner to fulfill the patrolling task. When the new task of opening a door is triggered, Digit navigates to the door and opens it by pushing the door with its left arm.

Localization of the quadcopter is performed with an overhead Vicon motion capture system (servo rate 20 Hz). The LTL framework computes setpoints for the quadcopter to track, and tracking is performed using a quadratic programming based controller, also operating at 20 Hz.

The experiment was set up to simulate a disaster relief scenario where the two robots collaborate in a disaster-struck house. Two desks (simulating walls) and a door are placed inside the Vicon motion capture area, with the desks as static obstacles and the door as the resolvable obstacle (i.e., can be opened). The uncertain region is represented using a tray with bricks (boxed in green in Fig. 5). Hence, Digit is tasked with patrolling between $s_{\text{uncertain}} = s_{43}$ and s_{awayC} in the center room.

During the hardware experiment, new goals are extracted from the automaton synthesized offline. Since commands extracted from the automaton take negligible time, the FSM
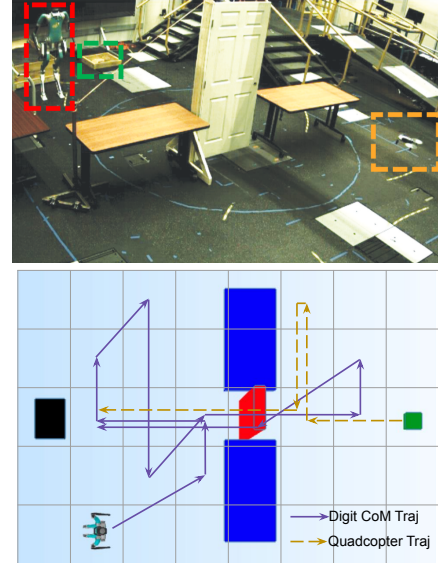

Fig. 5: Hardware experiment Setup. In the top figure, the desks represent walls — static obstacles in the LTL setup, and the door represents the resolvable obstacle. The tray boxed in green represents the uncertain region. The robot boxed in red is the Digit bipedal robot and the one boxed in blue is the Bebop quadcopter at initial positions. In the bottom figure, the robots' motion trajectories are represented via arrowed lines in a similar simulated environment.

can generate control policies at runtime. The hardware experiments showed that the proposed LTL framework was able to command the robot team to successfully open the door, sense the uncertain region, and fulfill the original tasks. The experiment was recorded in the video attachment to this submission.

The hardware implementation reveals a few new challenges. First, the synchronization between the two different robots makes the collaborative task less efficient as they reach the setpoints at different times, which results in longer time for the task completion. The initial LTL synthesis took around one minute, leaving Digit standing, and the drone airborne and exhausting their battery. Efforts can be made to reduce the initial synthesis time. Additionally, incorporating a high-performance whole-body motion controller for Digit such as [27] would enable improved walking performance, though that is beyond the scope of this work.

VIII. CONCLUSION AND FUTURE WORK

We presented a generalizable approach to identifying and resolving environment assumption violations discovered at runtime by automatically leveraging the capability of heterogeneous agents. This allows the team of agents to recover from cases where their objectives become unrealizable due to runtime-observed violations. We implemented this approach in a grid world simulation and generated safe 3D CoM motion plans for a bipedal robot and a quadcopter. Future directions of research include implementing deadlock resolution strategies and further developing the non-resynthesis solution to enable more complex runtime-assignable behaviors. Additionally, more complex sets of robot teams can be constructed to resolve a wider range of obstacles using the proposed approach. Future work can be done to catalogue a comprehensive library of robots, capabilities, and obstacle resolutions to enable broader applications, such as multi-quadcopter teaming to deliver a battery for Digit charging and executing long-duration navigation for package delivery. Robots that can pull doors open and remove debris can be added to the robot team to form a versatile team that aims for search and rescue applications. Migrating towards signal temporal logic is also of future research interest so that specifications about resolution time can be included, enabling the proposed approach to be deployed in timing-critical missions.

REFERENCES

[1] Y. Emam, S. Mayya, G. Notomista, A. Bohannon, and M. Egerstedt, "Adaptive task allocation for heterogeneous multi-robot teams with evolving and unknown robot capabilities," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 7719–7725.

[2] J. L. Kit, A. G. Dharmawan, D. Mateo, *et al.*, "Decentralized multi-floor exploration by a swarm of miniature robots teaming with wall-climbing units," in *International Symposium on Multi-Robot and Multi-Agent Systems*, 2019, pp. 195–201.

[3] Z. Zhou, D. J. Lee, Y. Yoshinaga, S. Balakirsky, D. Guo, and Y. Zhao, *Reactive task allocation and planning for quadrupedal and wheeled robots*, 2021.

[4] T. Wongpiromsarn, U. Topcu, and R. Murray, "Synthesis of control protocols for autonomous systems," *Unmanned Systems*, vol. 01, pp. 21–39, Jul. 2013.

[5] R. Tedrake, *Drake: Model-based design and verification for robotics*, 2019. [Online]. Available: https://drake.mit.edu.

[6] N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reactive(1) designs," in *Verification, Model Checking, and Abstract Interpretation*, Springer, 2006, pp. 364–380.

[7] J. Liu, N. Ozay, U. Topcu, and R. M. Murray, "Synthesis of reactive switching protocols from temporal logic specifications," *IEEE Transactions on Automatic Control*, vol. 58, no. 7, pp. 1771–1785, 2013.

[8] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.

[9] Z. Gu, N. Boyd, and Y. Zhao, "Reactive locomotion decision-making and robust motion planning for real-time perturbation recovery," in *IEEE International Conference on Robotics and Automation*, 2022.

[10] A. Pnueli and R. Rosner, "On the synthesis of a reactive module," in *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Austin, TX, USA: Association for Computing Machinery, 1989, pp. 179–190.

[11] K. W. Wong, R. Ehlers, and H. Kress-Gazit, "Correct high-level robot behavior in environments with unexpected events," in *Robotics: Science and Systems*, 2014.

[12] S. C. Livingston, R. M. Murray, and J. W. Burdick, "Backtracking temporal logic synthesis for uncertain environments," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 5163–5170.

[13] W. Li, L. Dworkin, and S. A. Seshia, "Mining assumptions for synthesis," in *ACM/IEEE International Conference on Formal Methods and Models for Codesign*, IEEE, 2011, pp. 43–50.

[14] R. Ehlers and U. Topcu, "Resilience to intermittent assumption violations in reactive synthesis," in *International Conference on Hybrid Systems: Computation and Control*, 2014, pp. 203–212.

[15] R. Majumdar, E. Render, and P. Tabuada, "Robust discrete synthesis against unspecified disturbances," in *International conference on Hybrid systems: computation and control*, 2011, pp. 211–220.

[16] V. Raman and H. Kress-Gazit, "Automated feedback for unachievable high-level robot behaviors," in *2012 IEEE International Conference on Robotics and Automation*, 2012, pp. 5156–5162.

[17] K. W. Wong, R. Ehlers, and H. Kress-Gazit, "Resilient, provably-correct, and high-level robot behaviors," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 936–952, 2018.

[18] J. van den Berg, M. Stilman, J. Kuffner, M. Lin, and D. Manocha, "Path planning among movable obstacles: A probabilistically complete approach," in *Algorithmic Foundation of Robotics VIII: Selected Contributions of the Eight International Workshop on the Algorithmic Foundations of Robotics*. Springer Berlin Heidelberg, 2010, pp. 599–614.

[19] A. Krontiris and K. Bekris, "Dealing with difficult instances of object rearrangement," in *Proceedings of Robotics: Science and Systems*, Rome, Italy, Jul. 2015.

[20] V. Vasilopoulos, Y. Kantaros, G. J. Pappas, and D. E. Koditschek, *Technical report: Reactive planning for mobile manipulation tasks in unexplored semantic environments*, 2021. arXiv: 2011.00642 [cs.RO].

[21] Agility Robotics, in *http://www.agilityrobotics.com/*.

[22] H. Kress-Gazit, T. Wongpiromsarn, and U. Topcu, "Correct, reactive, high-level robot control," *IEEE Robotics & Automation Magazine*, vol. 18, no. 3, pp. 65–74, 2011.

[23] R. Ehlers and V. Raman, *Slugs: Extensible gr(1) synthesis*, Springer, 2016.

[24] J. Warnke, A. Shamsah, Y. Li, and Y. Zhao, "Towards safe locomotion navigation in partially observable environments with uneven terrain," in *IEEE Conference on Decision and Control*, 2020, pp. 958–965.

[25] S. Bharadwaj, R. Dimitrova, and U. Topcu, "Synthesis of surveillance strategies via belief abstraction," in *IEEE Conference on Decision and Control*, IEEE, 2018, pp. 4159–4166.

[26] Y. Zhao, B. R. Fernandez, and L. Sentis, "Robust optimal planning and control of non-periodic bipedal locomotion with a centroidal momentum model," *The International Journal of Robotics Research*, vol. 36, no. 11, pp. 1211–1242, 2017.

[27] A. Shamsah, J. Warnke, Z. Gu, and Y. Zhao, *Integrated task and motion planning for safe legged navigation in partially observable environments*, 2021.