

# Behavior Tree Capabilities for Dynamic Multi-Robot Task Allocation with Heterogeneous Robot Teams

Georg Heppner<sup>1</sup>

David Oberacker<sup>1</sup>

Arne Roennau<sup>1</sup>

Rüdiger Dillmann<sup>1</sup>

**Abstract**—While individual robots are becoming increasingly capable, the complexity of expected missions increases exponentially in comparison. To cope with this complexity, heterogeneous teams of robots have become a significant research interest in recent years. Making effective use of the robots and their unique skills in a team is challenging. Dynamic runtime conditions often make static task allocations infeasible, requiring a dynamic, capability-aware allocation of tasks to team members. To this end, we propose and implement a system that allows a user to specify missions using Behavior Trees (BTs), which can then, at runtime, be dynamically allocated to the current robot team. The system allows to statically model an individual robot’s capabilities within our *ros.bt.py* BT framework. It offers a runtime auction system to dynamically allocate tasks to the most capable robot in the current team. The system leverages utility values and pre-conditions to ensure that the allocation improves the overall mission execution quality while preventing faulty assignments. To evaluate the system, we simulated a find-and-decontaminate mission with a team of three heterogeneous robots and analyzed the utilization and overall mission times as metrics. Our results show that our system can improve the overall effectiveness of a team while allowing for intuitive mission specification and flexibility in the team composition.

## I. INTRODUCTION

With the widespread adoption of increasingly complex robotic systems, the need for new methods of defining a system’s high-level behavior while simultaneously considering complex skill sets is becoming apparent. Specialized systems like walking robots or AI-assisted drones are now commercially available from multiple manufacturers. These systems are not only accessible but also increasingly versatile. They are being equipped with added manipulators or specialized 3D sensors. This enhancement makes the idea of teams comprising diverse robots a highly realistic proposition. However, coordinating very heterogeneous robots and skills while dealing with dynamic team compositions requires a robust and flexible system to model, coordinate, and exchange robot skills and behaviors. Therefore, implementing a decentralized system optimized to solve the *Multi-Robot Task Allocation Problem (MRTAP)* for heterogeneous robots is a worthwhile endeavor. Solving the MRTAP has been a topic in robotics for many years. Gerkey and Mataric [1] provide a classification of sub-problems to the MRTAP. In the following we will look primarily at the “Single Task -

<sup>1</sup> Department of Interactive Diagnosis and Service Systems (IDS), FZI Research Center for Information Technology, Haid-und-Neu-Straße 10–14, 76131 Karlsruhe, Germany.

The research leading to these results has received funding in the ROBDEKON II project under the grant agreement No. 13N16540 by the German Federal Ministry of Education and Research (BMBF).




Fig. 1. Top: Gazebo Simulation with a heterogeneous robot team, Middle: 2D Cave map with target objects (blue) exploration goals (red) and paths. All three robots employ this map for navigation with the ROS *move\_base* stack, Bottom: Behavior Tree segment used in the Explore Capability.

Multi Robot - Instantaneous Assignment (*ST-MR-IA*)” sub-problem.

Behavior Trees have emerged as modular and easy-to-use method for specifying robot behavior [2], [3]. Works by Ogren [4] and Marzinotto *et al.* [5]–[9] propose the foundational concepts for BTs in robotics. Early on, they were also proposed for multi-robot systems. Colledanchise *et al.* [10] outline how BTs can be utilized for controlling multi-robot setups. Jeong *et al.* [11] describe a centralized method for controlling multiple mobile robots using a singular BT. To a similar end, we also proposed a BT framework for controlling teams of heterogeneous robots called *ros-bt.py* [12]. While the previously mentioned approaches allow the remote execution of static, predefined sub-trees, no approach can handle heterogeneous or dynamic abilities in the robot team yet. Non Behavior Tree approaches to the MRTAP have focused more on methods for modeling robot capabilities. The ASyMTRe framework proposed in [13], [14] facilitates collaboration by automatically synthesizing

a data flow graph between sensors, processors, and actuators for a given team and set of tasks. In the CoMutaR framework [15] this principle is extended to account for resource limitations and dynamic changes during runtime, at the cost of a sizable modeling overhead. Guerin *et al.* [16] propose a robot-independent mission BT format in addition to a model for a robot’s abilities. The approach is focused on single robot industrial applications where tasks should be described independently of the robot executing them. Assigning tasks to teams with degrading abilities over time is covered by [17]. While the team composition is dynamic for this approach, it requires a static set of available capabilities. A similar problem is solved in [18], where capabilities and task suitability are modeled as vectors, allowing for a probabilistic approach to determine an optimal task allocation over the mission runtime.

The MRTAP has been explored more thoroughly for homogeneous robot teams. Older approaches like *MURDOCH* [19] and *TraderBots* [20]–[22] aim to group related tasks and allocate them using auctions. More recent approaches include [18], [23]–[26], which focus on task synergies and optimizing multiple metrics during allocation. Current approaches for MRTAP are capable of generating optimized allocations for static robot teams, often using a centralized authority. Yet task allocations for dynamic teams of robots with heterogeneous capabilities still have many challenges that need to be addressed. Similarly, the restriction to a centralized task allocation process introduces problems for mobile robots and spontaneous cooperation. Thus, our goal is to realize a decentralized multi-robot task allocation system specifically for teams of heterogeneous robots. This includes providing an abstraction for the heterogeneous skills of individual team members. Decentralization is realized by employing an auction-based approach that enables each robot to act as the auctioneer, voiding the need for a centralized authority. These approaches contribute an extension to the *ros\_bt.py* Behavior Tree (BT) framework by introducing *capabilities* as a skill abstraction that can dynamically be auctioned off to robots in the same team.

In section II, we present the theoretical model of our proposed capabilities. The implementation of the model as part of the open-source *ros\_bt.py* library is described in section III. Section IV details our experimental evaluation of the system using a simulated find-and-decontaminate mission. Finally, section V summarizes this work and provides an outlook for future works.

## II. CAPABILITIES IN BEHAVIOR TREES FOR MRTA

Our previously proposed Behavior Tree framework [12] implements common BT functionality such as flow control, ticking, and sub-trees to model individual capabilities. Additionally it can be used to design complete missions by enabling long running tasks and utility evaluation of sub-trees. Finally, the possibility to execute parts of the BT on a distant system enables multi-robot task allocation for a team of heterogeneous robots as the mission can be spread over a group of robots based on their individual skills. However, the

proposed method relies on *shovables*. This special decorator node explicitly marks sub-trees suitable for remote execution, assigning them at the startup of the mission tree. This limits the possibility of changing the assignment during runtime, restricts reactivity of the parameters and results, and limits the complexity of the sub-trees. We therefore extended our previous work by specifically adding *capabilities* to the model and enable a dynamic distribution via *auction* of these within the team. All following definitions build upon the concepts & definitions introduced in [12].

### A. Capabilities in Behavior Trees

*Capabilities* abstract individual robot skills as node representations of semantically meaningful operations, e.g., *MoveToPose* or *OpenDoor*, in a BT. They do this by providing functionality similar to a sub-tree, but are assigned at runtime. This enables dynamic evaluation, usage of a suitable implementation, distribution of its execution via auction, and handling input and output parameters across systems. Before ticking the BT, the actual sequence of BT nodes that perform the skill described by the capability is not known, which requires some special handling.

**Definition II.1** (Capabilities). A capability  $c$  is a Node in the BT that requests the execution of an action with a single semantic meaning without providing the required BT.

$$\forall c \in \mathcal{C} : c \in \mathcal{N}_T$$

Where  $\mathcal{N}_T$  denotes the set of all nodes in a BT and  $\mathcal{C}$  the set of all capabilities. Capabilities are executor leaf nodes.

**Definition II.2** (Parameters for Capabilities). A single Parameter  $p$  is given with the Triple  $p = (n, k, t)$  consisting of a node  $n$ , kind  $k$  (input, output, option) and data-type  $t$ .  $\mathbb{P}$  denote the set of all parameters,  $\mathcal{P}(\mathbb{P})$  its power set.

$$\begin{aligned} \text{inputs} : \mathcal{C} &\rightarrow \mathcal{P}(\mathbb{P}) & \text{outputs} : \mathcal{C} &\rightarrow \mathcal{P}(\mathbb{P}) \\ \text{inputs}(c) &= \{p = (n, k, t) \mid p \in \mathbb{P} \wedge n = c \wedge k = \text{input}\} \\ \text{outputs}(c) &= \{p = (n, k, t) \mid p \in \mathbb{P} \wedge n = c \wedge k = \text{output}\} \end{aligned}$$

Capabilities cannot have option values, as they are incompatible with the dynamic binding:

$$\forall c \in \mathcal{C} : \neg \exists p \in \mathbb{P} : p = (n, k, t) \mid n = c \wedge k = \text{option}$$

**Definition II.3** (Capability Implementation). A capability node only provides the model information of the functionality. To actually do something a capability implementation  $i$  is required. An implementation  $i$  contains a complete BT, given as BT-graph  $\mathcal{G}_T(\mathcal{N}_T, \mathcal{E}_T, \imath)$  where  $\mathcal{N}_T$  contains all of the nodes which might also be capabilities,  $\mathcal{E}_T$  defines the edges between them,  $\imath$  their order. Each implementation is specific to a single robot, specified by its world state  $\mathcal{W} \in \Sigma^x; x \in \mathbb{N}$ . The world state is build with an arbitrary data alphabet  $\Sigma$  which models the know information about the world for a single robot, including its sensor data and other states. The data-graph  $\mathcal{G}_D(\mathcal{N}_D, \mathcal{E}_D)$  models the data interactions with the parameters where  $\mathcal{N}_D$  are the data nodes and  $\mathcal{E}_D$  their edges.

$$i = (\mathcal{W}, \mathcal{G}_T(\mathcal{N}_T, \mathcal{E}_T, \imath), \mathcal{G}_D(\mathcal{N}_D, \mathcal{E}_D))$$

**Definition II.4** (Capability to Implementation connection). Robots can exchange parts of their world state with the function  $nearbyWorlds(\mathcal{W})$ , therefore  $\mathcal{I}$  denotes the set of all capability implementations over all robots currently in the same team. By explicitly marking a capability as local, only locally available implementations are considered. Each implementation  $i$  corresponds to a single capability  $c$ . The association between the two is defined by the  $c_{impl}$  function:

$$c_{impl} : \mathcal{I} \rightarrow \mathcal{C}$$

$c_{impl}$  is surjective but not injective, meaning every capability is associated with one up to infinity implementations. The  $impl$  function on the other hand returns all currently executable implementations for a given capability  $c$ :

$$\begin{aligned} impl : \mathcal{C} &\rightarrow \mathcal{P}(\mathcal{I}) \\ impl(c) &= \{i \mid i \in \mathcal{I} \wedge c_{impl}(i) = c \wedge \\ &\quad executable(c, i) \wedge validate_{precondition}(i)\} \end{aligned}$$

**Definition II.5** (Capability Execution). An implementation can be executed if it implements a capability, fulfills all the preconditions and has a valid execution location.

$$\begin{aligned} executable(c, i) &= (\mathcal{W}_i = \mathcal{W}_c) \vee \\ &\quad (\neg required_{local}(c) \wedge \\ &\quad \exists Env_T = (\mathcal{G}_T(\mathcal{N}_T, \mathcal{E}_T, \iota), -, -, \mathcal{W}_i) : \\ &\quad \exists n \in \mathcal{N}_T : n = RemoteCapabilitySlot) \end{aligned}$$

$\mathcal{W}_i$  denotes the world state where the implementation  $i$  is located,  $\mathcal{W}_c$  the world state where the capability  $c$  is located. This is either the case if the implementation is on the same robot as the interface or remote execution is allowed and the robot that has the implementation currently has a *RemoteCapabilitySlot*.

**Definition II.6** (Preconditions). To ensure that a given capability implementation  $i$  can be executed correctly during runtime preconditions can be specified. Preconditions can require a set of capability interfaces to be executed successfully before the implementation runs. For each required capability it can be specified if they should have been executed on the local robot or on any nearby robot.  $\mathcal{R}$  describes the set of all possible preconditions.

$$\mathcal{R} = \{(c, k) \mid c \in \mathcal{C}; k \in \{\text{local}, \text{remote}\}\}$$

The precondition function returns all preconditions for a given capability implementation.

$$precondition : \mathcal{I} \rightarrow \mathcal{P}(\mathcal{R})$$

Before  $i$  is executed every precondition is validated against the current runtime context:

$$\begin{aligned} validate_{precondition}(i) &= \forall r = (c, k) \in precondition(i) : \\ &\quad (k = \text{local} \wedge c \in \mathcal{N}_T) \vee \\ &\quad (k = \text{remote} \wedge \exists \mathcal{G}_T(\mathcal{N}'_T, \mathcal{E}'_T, \iota') \\ &\quad \in nearbyTrees(\mathcal{W}) : c \in \mathcal{N}'_T) \end{aligned}$$

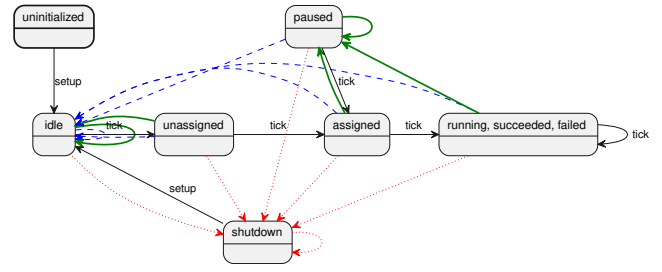


Fig. 2. State transitions for *Capability* BT nodes when applying actions. For clarity the error state/action has been left out. To further simplify: the red dotted edges are representing the *shutdown* action, blue dashed edges the *reset* action and the green solid lines the *untick* action. Lastly, the *running*, *succeeded*, and *failed* state haven been combined.

Where  $nearbyTrees$  is a function that returns the BT of nearby worlds. If a precondition cannot be validated the implementation execution will fail before ticking any node.

**Definition II.7** (Dynamic Binding). As the capabilities are evaluated at runtime (dynamic binding), the node State definition ([12]) needs to be extended to handle changing implementations during runtime. Two additional states are added:

$$\begin{aligned} state(n) \in \mathcal{N}_T = \{ &uninitialized, error, idle, \\ &succeeded, failed, running, \\ &shutdown, \mathbf{unassigned}, \mathbf{assigned} \} \end{aligned}$$

As the states can apply to all nodes, for example if a control flow node contains a capability, the behavior needs to be defined for all of them when using the update function  $update(n, a, Env_T)$  which determines the new node world state after applying an action  $a$ . For control-flow nodes, the new states act identical to the running state. Non-capability leaf nodes cannot reach the new states, thus they are not relevant for them. Capabilities follow the state transitions outlined in figure 2. An assignment function  $assign_c(c)$  allows retrieving the currently assigned implementation.

$$\begin{aligned} assign_c(c) : \mathcal{C} &\rightarrow \mathcal{I} \cup \{\emptyset\} \\ c \rightarrow \begin{cases} \emptyset & state(c) \in \{idle, shutdown, error, \\ & unassigned\} \\ i \in impl(c) & state(c) \in \{assigned, succeeded, \\ & running, paused, failed\} \end{cases} \end{aligned}$$

When returning an implementation the value out of the set  $impl(c)$  is determined in the *unassigned* state. While in the *assigned*, *paused*, *running*, *succeeded*, *failed* state, the value returned by  $impl(c)$  has to remain unchanged. Upon receiving a reset or shutdown action, the value will revert to  $\emptyset$ .

### B. Dynamic Capability Execution

The *shovables* we proposed in our previous work enable the remote execution of a sub-tree by extracting relevant parts of the  $Env_T$  and forwarding it to a *slot* selected during setup of the tree where the execution is handled before an updated

environment  $Env'_T$  is returned and integrated. Changes to parameters values and the selected *slot* are not supported during execution time, both a requirement for dynamic capability execution. To tackle this, and to ensure that updates are passed between the Environment  $Env_T$  of the capability and the environment  $Env^i_T$  of the implementation, which do not need to be the same, *CapabilityIOBridges* are introduced:

**Definition II.8** (Capability IO Bridge). For any given capability  $c$  there exists an  $c_{in}$  (Input Bridge) and an  $c_{out}$  (Output Bridge):

$$\begin{aligned} \forall c \in \mathcal{C} : \exists c_{in} \in \mathcal{C}_{in} : c_{in} &= br_{in}(c) \\ \forall c \in \mathcal{C} : \exists c_{out} \in \mathcal{C}_{out} : c_{out} &= br_{out}(c) \end{aligned}$$

$br_{in}(c)$  and  $br_{out}(c)$  are bijective functions associating bridges with capabilities.  $br_{in}(c_{br})$  and  $br_{out}(c_{br})$  are the respective inverse functions. Similarly to capabilities, both bridges are BT leaf nodes.

**Definition II.9** (Capability Input and Output Bridge). Capability input and output bridges pass the parameters from the capability to the implementation and back by acting as adapter. The input bridge nodes publishes the inputs passed to the capability as outputs, while the output bridge node similarly provides inputs that are passed to the outputs of the capability:

$$\begin{aligned} outputs(c_{in}) : \mathcal{C}_{in} &\rightarrow \mathcal{P}(\mathbb{P}) \\ c &\rightarrow \{p = (n, k, t) \mid p \in inputs(br_{in}(c_{in}))\} \\ inputs(c_{out}) : \mathcal{C}_{out} &\rightarrow \mathcal{P}(\mathbb{P}) \\ c &\rightarrow \{p = (n, k, t) \mid p \in outputs(br_{out}(c_{in}))\} \end{aligned}$$

Input bridges do not have inputs, while output bridges do not have outputs:

$$\begin{aligned} inputs(c_{in}) : \mathcal{C}_{in} &\rightarrow \mathcal{P}(\mathbb{P}) \\ outputs(c_{out}) : \mathcal{C}_{out} &\rightarrow \mathcal{P}(\mathbb{P}) \\ c &\rightarrow \emptyset \end{aligned}$$

While the state of the capability interface  $state(c) = running$ , during each call of  $Env'_T = update(c, tick, Env_T)$ , the following is ensured after every tick:

$$\begin{aligned} \forall c \in \mathcal{C} : c_{in} &= br_{in}(c) : \\ \forall p &= (n, k, t) \in inputs(c) : \\ \forall p' &= (n', k', t') \in outputs(c_{in}) : \\ n &= n' \\ \implies value(p, \mathcal{W}) &= value(p', \mathcal{W}'), \\ \forall c \in \mathcal{C} : c_{out} &= br_{out}(c) : \\ \forall p &= (n, k, t) \in outputs(c) : \\ \forall p' &= (n', k', t') \in inputs(c_{out}) : \\ n &= n' \\ \implies value(p, \mathcal{W}) &= value(p', \mathcal{W}'). \end{aligned}$$

Where  $value(p, W)$  is the parameter value function.  $W$  and  $W'$  are the respective world states of the capability interface and the implementation.

In order to execute the capabilities on arbitrary system the *RemoteCapabilitySlot* is introduced:

**Definition II.10** (Remote Capability Slot). A remote capability slot  $c_r$  is a Node that can execute local capability implementations  $i \in \mathcal{I}$  on request from remote robots. Only one implementation can be executed at a time. Upon receiving an execution request, containing the capability  $c$  and the requested implementation  $i = (\mathcal{W}, \mathcal{G}_T(\mathcal{N}_T, \mathcal{E}_T, \lambda), \mathcal{G}_D(\mathcal{N}_D, \mathcal{E}_D))$ , the node creates a new execution environment called  $Env^c_T$  with:

$$Env^c_T = (\mathcal{G}_T, \mathcal{G}_D, \Sigma, \mathcal{W})$$

This execution environment exists in the world state of the BT  $c_r$  is a part of while using an independent  $\mathcal{G}_T$  and  $\mathcal{G}_D$ . Actions for the new environment  $Env^c_T$  are passed through the  $c_r$  node, similar to a decorator node. Thus, the  $Env^c_T$  is ticked at the same rate as the environment of  $c_r$ . Parameter values within the  $Env^c_T$  are exchanged with the remote robots  $c$  by using the *CapabilityIOBridges* and the state of the root of  $\mathcal{G}_T$  is transmitted. When an implementation finishes its execution, all traces of it are removed from the environment the *RemoteCapabilitySlot* is part of, resetting it to the state before accepting the request. It uses a similar state machine as regular capability interfaces, depicted in figure 2. Main differences are that the unassigned, assigned and running state all revert back to the unassigned state after unticking and that the transition from unassigned to the assigned state requires an external signal originating from the capability that has requested the remote operation.

### C. Capability Behavior

To clarify the behavior of the capability, their respective behaviors during the relevant states are explained:

**Definition II.11** (Behavior in the Unassigned State). When  $Env'_T = update(c, tick, Env_T)$  is called while  $state(c) = unassigned$ , the capability will try to determine a concrete return value for the  $assign_c(c)$  function. The assignments will be determined by an auction-based system. In  $Env'_T$  the following holds:  $state(c) = assigned$  and  $assign_c(c) \neq \emptyset$ .

**Definition II.12** (Behavior in the Assigned State). When  $Env'_T = update(c, tick, Env_T)$  is called while  $state(c) = assigned$ , the capability interface will setup the tree specified by the  $assign_c(c)$  function. When a implementation on the local robot is selected for execution, it will be setup similarly to the sub-trees specified in [12] with the slot being the capability  $c$  itself. If the execution is on a remote robot, the implementation will be setup in an available *RemoteCapabilitySlot*. The created execution environment, either local or remote, is called  $Env^c_T$ , with  $r \in \mathcal{N}_T$  representing the root of the implementation BT. If this completes, the node will proceed to the running state.

**Definition II.13** (Behavior in the Running State). When  $Env'_T = update(c, tick, Env_T)$  is called while  $state(c) = running$ , the capability interface will forward the tick to the implementation execution environment. As described




Fig. 3. Schematic overview of the components required to implement the concept. Components marked in red, were implemented as part of this work.

later, before the tick, the current input values of  $inputs(c)$  will be passed into  $Env_T^c$  using the *CapabilityIOBridge*.  $Env_T^c = update(r, tick, Env_T^c)$  is called. The new values for  $outputs(c)$  are extracted from  $Env_T^c$ , using the *CapabilityIOBridge*, as well as the state  $state(r) \in S$ .  $state(c) = state(r)$  is set and the values for  $outputs(c)$  are updated in  $Env_T^c$ .

### III. IMPLEMENTATION

The presented concept was implemented as an extension to our *ros\_bt\_py* framework [12] (Figure 3). Each component is implemented as a Robot Operating System (ROS) node written in Python, whereas the interfaces are ROS services or ROS topics. The implementation of *ros\_bt\_py* was made public as open source project<sup>a</sup> and has recently been updated with the presented capability support.

*Capabilities*, *RemoteCapabilitySlots*, *IO inputs* and *IO outputs* are implemented as abstract BT nodes, which can be instantiated for specific input and output configurations. *Capability Implementations* are realized as sub-trees, where only IO input and output nodes are used to pass information across the sub-tree boundaries.

The **Capability Repository** implements a distributed database of all available capabilities within the current team. Additionally, it manages the available capability implementations on the local robot. Both capabilities and implementations can be added and removed during runtime.

Local and remote execution management for capabilities is realized in the **Mission Control** component. For local execution, it selects the best local implementation from the capability repository based on calculated utility values [12]. To realize remote execution, the best local implementation, and the corresponding capability are passed to the **Assignment System**. Upon receiving the best available remote robot, the local mission control will trigger the execution by communicating with the remote mission control. During local or remote execution, it will receive heartbeat messages with reports on execution progress.

The aforementioned **Assignment System** is an abstract component for determining the best available team member to execute a capability. In the reference implementation, the component is realized as an open single-item auction

<sup>a</sup>[https://github.com/fzi-forschungszentrum-informatik/ros\\_bt\\_py](https://github.com/fzi-forschungszentrum-informatik/ros_bt_py)




Fig. 4. Simplified Behavior Tree describing the overall mission. Each Action node is a *Capability* that is assigned to a robot for execution. The root node is a **parallel** node that ticks all its children in parallel.

system with re-auctions during execution time. During the bid period, the weighted utility values from all eligible robots in the team are collected. Robots are greedy bidders running as many tasks as possible while minimizing costs. Bids are sent to auctioneer event-based with multiple triggers:

- Request for Bids at the start of the auction.
- Team composition changes.
- If the current auction winner's costs exceed 120% of his winning bid.

To enable decentralized coordination of the team, each robot executes all components (Capability Repository, Mission Control, Assignment System) as part of his framework instance.

### IV. EVALUATION

To evaluate our previously presented approach, we simulate a heterogeneous robot team performing a find-and-decontaminate mission in a cave environment. The goal is for the system to automatically assign the different tasks within the mission to capable robots while adapting the plan to current runtime conditions. Task assignment is based on the position of the robot and the robots ability to perform the task. Robot failures, team expansions, and uncertainty of successful task completion were induced to evaluate the system's ability to adapt during runtime.

#### A. Mission

The primary mission is to find contaminated objects in the cave and move them to the decontamination area. Thus, there are three tasks for the robots to execute: **Explore** reveals nearby objects and further exploration points, **Identify** checks if a found object is contaminated, and **Decontaminate** moves contaminated objects to the decontamination area. As the aim of the evaluation is the task allocation, we simplified the task execution by using a *simulation manager*. The simulation manager offers ROS services to simulate the sensor results and manipulation operations during task execution. To make the simulated execution more realistic, the simulation manager waits for randomized intervals before returning results, with a randomized chance for each operation to fail. In figure 1, the tasks 2D-Poses within the map can be seen. Figure 4 shows the simplified BT encoding the overall mission behavior. Found object poses during exploration are exchanged using a globally shared database.




Fig. 5. Task allocation during a mission run with a single Husky robot. Arrows denote time points when a task, with the ID as seen in figure 1, was completed.

### B. Simulation Environment

Gazebo<sup>b</sup> with ROS 1 Noetic is used as the framework for the simulation. The underground cave environment is based on the environment presented in [27]. Our robot team comprises a subset of three heterogeneous robots: Two Clearpath Husky, and a Parrot Bebop.

TABLE I

OVERVIEW OF THE CAPABILITIES OF EACH TEAM MEMBER. THE CHECK MARK INDICATES THAT A ROBOT IMPLEMENTS THE CAPABILITY. NUMBERS IN BRACKETS INDICATE THE RUNTIME AND EXPLORED RADIUS.

Robot Name	Distance Cost Factor	Explore	Identify	Decontaminate
<b>Husky</b>	2.0	✓ (60 sec, 15m)	✓ (60 sec)	✓ (8* sec)
<b>Bebop</b>	1.0	✓ (10 sec, 30m)	✗ (- sec)	✗ (- sec)

Table I shows the cost estimation parameters for each robot and which capabilities they provide. All robots start within the start area, as seen in figure 1, and use the same BT consisting only of a RemoteCapabilitySlot.

### C. Dynamic Assignment

As seen in Figure 4, the capabilities abstract robot skills as an easy-to-understand and reusable module that enables quick mission building. During each tick of the capabilities, the best executor is evaluated, dynamically assigning the part of the mission to the most eligible robot. Figure 5 shows the task allocation for a mission run with a single Husky robot. As per the robot's bidding strategy, the auction-based task assignment preferred the execution of tasks with lower overall costs. As the distance to the task location is the major cost factor, this expresses itself in the execution of the physically closest task. Figure 6 depicts a mission run starting with two Husky robots. During this run, a Bebop drone joins for a certain period, and a Husky fails for a short period. Here, dynamic task allocation based on each robot's reported costs can be observed. The explore task is allocated to the Bebop drone immediately after it joins the team, as it reports lower costs even though it has to travel a longer




Fig. 6. Task allocation over the mission runtime when using a full team consisting of two Huskys and a Bebop. Arrows denote time points when a task, with the ID as seen in figure 1, was completed. Green arrows indicate a team member joining the team, whereas red arrows indicate a team member leaving.

distance initially. The existing allocation is adapted through re-auctions to minimize the overall costs and leverage the possible task parallelism from employing a robot team.

At all times, the system provided a valid, but due to the simple auction strategy not always globally optimal, assignment of tasks to robots based on their capabilities.

## V. CONCLUSIONS AND FUTURE WORKS

In this paper, we have presented Behavior Tree Capabilities as an extension of the previously published `ros_bt_capability` to conveniently structure robot skills and dynamically distribute them within a team of heterogeneous robots. We presented the theoretical concept that comes with the notion of preconditions, IOBridges, and RemoteCapabilitySlots and how it is implemented. The proposed system was evaluated with a simulated mission, showcasing the ability to distribute tasks in the team dynamically and to optimize the overall mission by selecting the most suitable systems. We have shown that the Behavior Tree Capabilities are an intuitive and versatile system to abstract individual robot behavior and can be used to enable spontaneous cooperation within a team. By using capabilities as a model in the BT mission definition, while allowing for runtime selection of the respective robot-dependent implementations, tasks can be handled robot-specific while remaining modular. Such flexibility greatly enhances the option of using multiple heterogeneous robots and spontaneously adding new systems to the mission. By building the assignment system into the core library and defining an implementation mapping as a core principle, we enable an ad-hoc distribution of the capabilities to the most suitable system and, therefore, a decentralized solution to the MRTAP. By modeling preconditions and allowing (re-)distribution, we reduce the required apriori assumptions and knowledge about a mission. Instead the focus shifts on the development of robust, reusable skills for individual robots. Currently, we are porting the library to ROS 2 and are considering the transfer and translation of capabilities as a form of skill learning in the future.

<sup>b</sup>[http://wiki.ros.org/gazebo\\_ros\\_pkgs](http://wiki.ros.org/gazebo_ros_pkgs)

## REFERENCES

- [1] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004-09-01. DOI: 10.1177/027836490404045564.
- [2] M. Iovino, E. Scukins, J. Styurd, P. Ögren, and C. Smith, "A survey of behavior trees in robotics and ai," *Robotics and Autonomous Systems*, vol. 154, p. 104096, 2022, ISSN: 0921-8890. DOI: 10.1016/j.robot.2022.104096.
- [3] R. Ghzouli, T. Berger, E. B. Johnsen, S. Dragule, and A. Wkasowski, "Behavior trees in action: A study of robotics applications," in *Proceedings of the 13th ACM SIGPLAN International Conference on Software Language Engineering*, ser. SLE 2020, Virtual, USA: Association for Computing Machinery, 2020, pp. 196–209. DOI: 10.1145/3426425.3426942.
- [4] P. Ögren, "Increasing modularity of UAV control systems using computer game behavior trees," in *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, 2012-06-09. DOI: 10.2514/6.2012-4458.
- [5] A. Marzinotto, M. Colledanchise, C. Smith, and P. Ögren, "Towards a unified behavior trees framework for robot control," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014-05, pp. 5420–5427. DOI: 10.1109/ICRA.2014.6907656.
- [6] M. Colledanchise, R. M. Murray, and P. Ögren, "Synthesis of correct-by-construction behavior trees," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017-09, pp. 6039–6046. DOI: 10.1109/IROS.2017.8206502.
- [7] M. Colledanchise and L. Natale, "On the implementation of behavior trees in robotics," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5929–5936, 2021-07, Conference Name: IEEE Robotics and Automation Letters. DOI: 10.1109/LRA.2021.3087442.
- [8] M. Colledanchise and P. Ögren, *Behavior Trees in Robotics and AI: An Introduction*. Boca Raton: CRC Press, 2018-07-26, 208 pp. DOI: 10.1201/9780429489105.
- [9] M. Colledanchise and P. Ögren, "How Behavior Trees Modularize Hybrid Control Systems and Generalize Sequential Behavior Compositions, the Subsumption Architecture, and Decision Trees," *IEEE Transactions on Robotics*, vol. 33, no. 2, pp. 372–389, 2017-04. DOI: 10.1109/TRO.2016.2633567.
- [10] M. Colledanchise, A. Marzinotto, D. V. Dimarogonas, and P. Ögren, "The advantages of using behavior trees in multi-robot systems," in *Proceedings of ISRA 2016: 47th International Symposium on Robotics*, 2016-06, pp. 1–8.
- [11] S. Jeong, T. Ga, I. Jeong, and J. Choi, "Behavior tree-based task planning for multiple mobile robots using a data distribution service," in *2022 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, Sapporo, Japan: IEEE Press, 2022, pp. 1791–1798. DOI: 10.1109/AIM52237.2022.9863364.
- [12] G. Heppner, N. Berg, D. Oberacker, N. Spielbauer, A. Roennau, and R. Dillmann, "Distributed behavior trees for heterogeneous robot teams," in *Proceedings of CASE 2023: 19th International Conference on Automation Science and Engineering*, Proceedings not yet publicized; Preprint available on arXiv, 2023-08.
- [13] F. Tang and L. Parker, "ASyMTRe: Automated synthesis of multi-robot task solutions through software reconfiguration," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2005-04, pp. 1501–1508. DOI: 10.1109/ROBOT.2005.1570327.
- [14] Y. Zhang and L. E. Parker, "IQ-ASyMTRe: Forming executable coalitions for tightly coupled multirobot tasks," *IEEE Transactions on Robotics*, vol. 29, no. 2, pp. 400–416, 2013-04. DOI: 10.1109/TRO.2012.2228135.
- [15] P. M. Shiroma and M. F. M. Campos, "CoMutaR: A framework for multi-robot coordination and task allocation," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009-10, pp. 4817–4824. DOI: 10.1109/IROS.2009.5354166.
- [16] K. R. Guerin, C. Lea, C. Paxton, and G. D. Hager, "A framework for end-user instruction of a robot assistant for manufacturing," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015-05, pp. 6167–6174. DOI: 10.1109/ICRA.2015.7140065.
- [17] M. Irfan and A. Farooq, "Auction-based task allocation scheme for dynamic coalition formations in limited robotic swarms with heterogeneous capabilities," in *2016 International Conference on Intelligent Systems Engineering (ICISE)*, 2016-01, pp. 210–215. DOI: 10.1109/INTELSE.2016.7475122.
- [18] B. Fu, W. Smith, D. Rizzo, M. Castanier, M. Ghaffari, and K. Barton, *Robust task scheduling for heterogeneous robot teams under capability uncertainty*, 2022-06-22. DOI: 10.48550/arXiv.2106.12111.
- [19] B. Gerkey and M. Mataric, "Sold!: Auction methods for multirobot coordination," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 758–768, 2002-10. DOI: 10.1109/TRA.2002.803462.
- [20] R. Zlot, A. Stentz, M. Dias, and S. Thayer, "Multi-robot exploration controlled by a market economy," in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, vol. 3, 2002-05, 3016–3023 vol.3. DOI: 10.1109/ROBOT.2002.1013690.

- [21] M. B. Dias, M. Zinck, R. Zlot, and A. Stentz, "Robust multirobot coordination in dynamic environments," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*, vol. 4, 2004-04, 3435–3442 Vol.4. DOI: 10.1109/ROBOT.2004.1308785.
- [22] M. B. Dias, "TraderBots: A new paradigm for robust and efficient multirobot coordination in dynamic environments," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, USA, 2004-01.
- [23] K. Rajchandar, R. Baskaran, P. K. Padmanabhan, and M. Rajmohan, "A novel fuzzy and reverse auction-based algorithm for task allocation with optimal path cost in multi-robot systems," *Concurrency and Computation: Practice and Experience*, vol. 34, no. 5, e6716, 2022, ISSN: 1532-0634. DOI: 10.1002/cpe.6716.
- [24] M. Alshaboti and U. Baroudi, "Multi-robot task allocation system: Fuzzy auction-based and adaptive multi-threshold approaches," *SN Computer Science*, vol. 2, no. 2, p. 87, 2021-02-10. DOI: 10.1007/s42979-021-00479-x.
- [25] X. Wen and Z.-G. Zhao, "Multi-robot task allocation based on combinatorial auction," in *2021 9th International Conference on Control, Mechatronics and Automation (ICCMA)*, 2021-11, pp. 27–32. DOI: 10.1109/ICCMA54375.2021.9646189.
- [26] T. Eijyne, G. Rishwaraj, and S. G. Ponnambalam, "Development of a task-oriented, auction-based task allocation framework for a heterogeneous multirobot system," *Sādhanā*, vol. 45, no. 1, p. 115, 2020-05-09. DOI: 10.1007/s12046-020-01330-4.
- [27] A. Koval, C. Kanellakis, E. Vidmark, J. Haluska, and G. Nikolakopoulos, *A subterranean virtual cave world for gazebo based on the DARPA SubT challenge*, 2020-04-21. DOI: 10.48550/arXiv.2004.08452.