

Adaptive Fault Tolerant Execution of Multi-Robot Missions using Behavior Trees

Michele Colledanchise,[†] Alejandro Marzinotto,[†] Dimos V. Dimarogonas,[‡] and Petter Ögren[†]

Abstract—Multi-robot teams offer possibilities of improved performance and fault tolerance, compared to single robot solutions. In this paper, we show how to realize those possibilities when starting from a single robot system controlled by a Behavior Tree (BT). By extending the *single robot BT* to a *multi-robot BT*, we are able to combine the fault tolerant properties of the BT, in terms of built-in fallbacks, with the fault tolerance inherent in multi-robot approaches, in terms of a faulty robot being replaced by another one. Furthermore, we improve performance by identifying and taking advantage of the opportunities of parallel task execution, that are present in the single robot BT. Analyzing the proposed approach, we present results regarding how mission performance is affected by *minor faults* (a robot losing one capability) as well as *major faults* (a robot losing all its capabilities). Finally, a detailed example is provided to illustrate the approach.

I. INTRODUCTION

Imagine a robot that is designed to perform maintenance on a given machine. The robot can open and close the cover, do fault detection and replace broken hardware (HW) components. However, this robot is fairly complex and easily breaks down. Thus, it is desirable to replace this big versatile robot with a team of smaller specialized robots, as shown in Figure 1. This paper shows how to modify the single robot Behavior Tree (BT) controlling the original robot, to a multi-robot BT, to be run in parallel on all new robots, thereby improving both fault tolerance and performance of the original BT.

The fault tolerance of the single robot BT, in terms of *fallbacks*, is improved by also adding the tolerance afforded by the redundancy achieved through having multiple robots, and the performance of the single robot BT is improved by executing the right tasks in parallel. In detail, the BT includes both so-called *sequence* and *fallback* compositions, as explained below. Some tasks that were earlier done in sequence, such as diagnosing different HW components, are now done in parallel, while other tasks that were also done in sequence, such as opening the cover and then diagnosing HW, are still done in a sequence. Similarly, some tasks that were earlier done as fallbacks, such as searching in storage 2 if no tools were found in storage 1, are now done in parallel, while other fallbacks, such as using the small screwdriver if the large one does not fit, are not done in parallel. Thus fault tolerance as well as performance is improved.

[†]The author is with the Centre for Autonomous Systems, Computer Vision and Active Perception Lab, School of Computer Science and Communication, The Royal Institute of Technology - KTH, Stockholm, Sweden.

[‡]The author is with the Centre for Autonomous Systems, Automatic Control Lab, School of Electrical Engineering, The Royal Institute of Technology - KTH, Stockholm, Sweden.




Fig. 1. Snapshot of a multi-robot mission execution. Two robots are replacing the broken parts, HW2 and HW4, (bottom) of a machine. One robot is out of order (top left), and two robots are resting (top right).

The contribution of this paper is that we show how fault tolerance and performance can be improved in a single robot BT, by adding more robots to the team and extending the BT into a multi-robot BT, using two important modifications. First we add a task assignment functionality, and then we identify and adapt the sequences and fall back compositions of the BT that are possible to execute in parallel.

The outline of this paper is as follows. First, in Section II we review related work. Then, in Section III we give an overview of the classical formulation of BTs. The main problem is stated in Section IV and the proposed solution is given in Section V. An extensive example illustrating the proposed approach is presented in Section VI, and we conclude the paper in Section VII.

II. RELATED WORK

BTs are a recent alternative to Controlled Hybrid Systems (CHSs) for reactive fault tolerant execution of robot tasks, and they were first introduced in the computer gaming industry [1]–[3], to meet their needs of modularity, flexibility and reusability in artificial intelligence for in-game opponents. Their popularity lies in its ease of understanding, its recursive structure, and its usability, creating a growing attention in academia [4]–[11]. In most cases, CHSs have memoryless transitions, i.e. there is no information where the transition took place from, a so-called *one way control transfer*. In BTs the equivalent of state transitions are governed by calls and return values being sent by parent/children in the tree structure, this information passing is called a *two way control transfers*. In programming languages, the replacement of a one way (e.g., GOTO statement) with a two way control transfer (i.e., Function Calls) made an improvement in readability and reusability [12]. Thus, BTs exhibit similar advantages as gained moving from GOTO to Function Calls in programming in the 1980s. Note however,

we do *not* claim that BTs are better than CHSs from a purely theoretical standpoint. Instead, the main advantage of using BTs lies in its ease of use, in terms of modularity and reusability, [8]. BTs were first used in [1], [2], in high profile computer games, such as the HALO series. Later work merged machine learning techniques with BTs’ logic [4], [5], making them more flexible in terms of parameter passing [6]. The advantage of BTs as compared to a general Finite State Machines (FSMs) was also the reason for extending the JADE agent Behavior Model with BTs in [7], and the benefits of using BTs to control complex missions of Unmanned Aerial Vehicles (UAVs) was described in [8]. In [10] the modular structure of BTs addressed the formal verification of mission plans.

In this work, we show how a plan for a multi-robot system can be addressed in a BT fashion, gaining all the advantages aforementioned in addition to the scalability that distinguishes a general multiagent system. Many existing works [13]–[15] stressed the problem of defining local tasks to achieve a global specification emphasizing the advantages of having a team of robots working towards a global goal. Moreover, [16], [17] introduce the concept of task delegation among agents in a multi-agent system, dividing task specification in *closed delegation*, where goal and plan are predefined, and *open delegation* where either only the goal is specified while the plan can be chosen by the agent, or the specified plan describes abstractly what actions have to be taken, giving to the agent some freedom in terms of how to perform the delegated task. In [18], it is shown how verifying the truth of preconditions on single agents becomes equivalent to checking the fulfillment of a global robot network through recursive calls, using a tree structure called *Task Specification Trees*.

Recent works present some advantages of implementing BTs in robotics applications [8], [19] making comparisons with the CHSs, highlighting their modularity and reusability. In [9], BTs were used to perform autonomous robotic grasping. In particular, it was shown how BTs enabled the easy composition of primitive actions into complex and robust manipulation programs. Finally, in [20] performance measures, such as success/failure probabilities and execution times, were estimated for plans using BTs. However, BTs are mostly used to design single agent behavior and, to the best of our knowledge, there is no rigorous framework in academia using the classical formulation of BTs for multi-robot systems.

III. BACKGROUND: FORMULATION OF BTs

Here, we briefly introduce an overview of BTs, the reader can find a detailed description in [8].

A BT is defined as a directed rooted tree where nodes are grouped into control flow nodes, execution nodes, and a root node. In a pair of connected nodes we call the outgoing node *parent* and the incoming node *child*. Then, the root node has no parents and only one child, the control flow nodes have one parent and at least one child, and the execution nodes are the leaves of the tree (i.e. they have no children and one

parent). Graphically, the children of a control flow node are sorted from its bottom left to its bottom right, as depicted in Figures 2-4. The execution of a BT starts from the root node. It sends *ticks*¹ to its child. When a generic node in a BT receives a tick from its parent, its execution starts and it returns to its parent a status *running* if its execution has not finished yet, *success* if its execution is accomplished (i.e. the execution ends without failures), or *failure* otherwise. Here we draw a distinction between three types of control flow nodes (selector, sequence, and parallel) and between two types of execution nodes (action and condition). Their execution is explained below.

Selector (also known as Fallback): When the execution of a selector node starts (i.e. the node receives a tick from its parent), then the node’s children are executed in succession from left to right, until a child returning success or running is found. Then this message is returned to the parent of the selector. It returns failure only when all the children return a status failure. The purpose of the selector node is to robustly carry out a task that can be performed using several different approaches (e.g. a motion tracking task can be made using either a 3D camera or a 2D camera) by performing each of them in succession until one succeeds. The graphical representation of a selector node is a box with a “?”, as in Fig. 2.

A finite number of BTs $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N$ can be composed into a more complex BT, with them as children, using the selector composition: $\mathcal{T}_0 = \text{Selector}(\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N)$.




Fig. 2. Graphical representation of a selector node with N children.

Sequence: When the execution of a sequence node starts, then the node’s children are executed in succession from left to right, returning to its parent a status failure (running) as soon as the a child that returns failure (running) is found. It returns success only when all the children return success. The purpose of the sequence node is to carry out the tasks that are defined by a strict sequence of sub-tasks, in which all have to succeed (e.g. a mobile robot that has to move to a region “A” and then to a region “B”). The graphical representation of a sequence node is a box with a “→”, as in Fig. 3.

A finite number of BTs $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N$ can be composed into a more complex BT, with them as children, using the sequence composition: $\mathcal{T}_0 = \text{Sequence}(\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N)$.

Parallel: When the execution of a parallel node starts, then the node’s children are executed in succession from left to right without waiting for a return status from any child before ticking the next one. It returns success if a given number of children $M \in \mathbb{N}$ return success, it returns

¹A tick is a signal that enables the execution of a child.




Fig. 3. Graphical representation of a sequence node with N children.

failure when the children that return running and success are not enough to reach the given number, even if they would all return success. It returns running otherwise. The purpose of the parallel node is to model those tasks separable in independent sub-tasks performing non conflicting actions (e.g. a multi object tracking can be performed using several cameras). The parallel node is graphically represented by a box with “ \Rightarrow ” with the number M on top left, as in Fig. 4. A finite number of BTs $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N$ can be composed into a more complex BT, with them as children, using the parallel composition: $\mathcal{T}_0 = \text{Parallel}(\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N, M)$.




Fig. 4. Graphical representation of a parallel node with N children.




Fig. 5. Graphical representation of action and condition nodes.

Action: When an action node starts its execution, then it returns success if the action is completed and failure if the action cannot be completed. Otherwise it returns running. The action node is represented in Fig. 5(a)

Condition: The condition node checks if a condition is satisfied or not. The return status is success or failure accordingly and it is never running. The condition node is represented in Fig. 5(b).

Root: The root node is the node that generates ticks. It is graphically represented by a box labeled with “ \emptyset ”.

IV. PROBLEM FORMULATION

The problem considered in this work is to define identical decentralized local BT controllers for the single robots of a multi-robot system to achieve a global goal. We first give some assumptions used throughout the paper, we then state the main problem.

Let \mathcal{S} be a set of symbols describing atomic actions (e.g. perform grasp, go to position) and let $\mathcal{R} \in \mathbb{N}$ be the set of robots in a multi-robot system, each of which can perform a finite collection of local tasks $\mathcal{L}_i \subseteq \mathcal{S}$, $i \in \mathcal{R}$ and $n = |\mathcal{R}|$ and let $\tilde{\mathcal{L}} = \bigcup_{i \in \mathcal{R}} \mathcal{L}_i$ be the set of all the local tasks. Now let $\mathcal{G} = \{g_1, g_2, \dots, g_v\}$ be a finite collection of global tasks executable by the multi-robot system and P_k be the set of global tasks running in parallel with g_k . We define $\psi(g_k)$ as

the finite set of local tasks which have to return success for the global task g_k to succeed.

Finally, let $\nu : \mathcal{G} \times \tilde{\mathcal{L}} \rightarrow \mathbb{N}$ and $\mu : \mathcal{G} \times \tilde{\mathcal{L}} \rightarrow \mathbb{N}$ respectively be functions that give the minimum number of robots needed, and the maximum number of robots assignable, to perform a local task, in order to accomplish a global task. Given these, we state the following assumptions:

Assumption 1: The minimum number of robots required to execute global tasks running in parallel does not exceed the total number of robots in the system: $\sum_{g_k \in P_h} \sum_{l_j \in \psi(g_k)} \nu(g_k, l_j) \leq n \forall g_h \in \mathcal{G}$.

Assumption 2: Each global task in \mathcal{G} can be performed by assigning to some robots in \mathcal{R} one of their local tasks: $\psi(g_k) \subseteq 2^{\tilde{\mathcal{L}}} \forall g_k \in \mathcal{G}$.

Problem 1: Given a multi-robot system defined as above and a global goal that satisfies Assumptions 1-2, design a local BT controller for each single robot to achieve this goal.

V. PROPOSED SOLUTION

In this section, we will first give an informal description of the proposed solution, involving the three subtrees \mathcal{T}_G (global tasks), \mathcal{T}_A (task assignment) and \mathcal{T}_{L_i} (local tasks). Then, we describe the design of these subtrees in detail. To illustrate the approach we give a brief example, followed by a description of how to create a multi agent BT from a single agent BT. Finally, we analyze the fault tolerance of the proposed approach.

The multi agent BT is composed of three subtrees running in parallel, returning success only if all three subtrees return success.

$$\mathcal{T}_i = \text{Parallel}(\mathcal{T}_A, \mathcal{T}_G, \mathcal{T}_{L_i}, 3) \quad (1)$$

Remark 1: Note that all robots run identical copies of this tree, including computing assignments of all robots in the team, but each robot only executes the task that they themselves are assigned to, according to their own computation.

The first subtree, \mathcal{T}_A is doing task assignment. Parts of the performance and fault tolerance of the proposed approach comes from this feature, making sure that a broken robot is replaced and that robots are assigned to the tasks they do best.

The second subtree, \mathcal{T}_G takes care of the overall mission progression. This tree includes information on task that needs to be done in sequence (using sequence composition), tasks that could either be done in sequence or in parallel (using parallel $M = N$ compositions), fallback tasks that should be tried in sequence (using selector compositions), and fallbacks that can be tried in parallel (using parallel $M = 1$ compositions). The execution of \mathcal{T}_G provides information to the assignment in \mathcal{T}_A . The fault tolerance of the proposed approach is improved by the fallbacks encoded in \mathcal{T}_G , and the performance is improved by parallelizing actions whenever possible, as described above.

The third subtree, \mathcal{T}_{L_i} uses output from the task assignment to actually execute the proper actions. We will now define the three subtrees in more detail.

a) *Definition of \mathcal{T}_A* : This tree is responsible for the reactive optimal task assignment, deploying robots to different local tasks according to the required scenario. The reactivity lies in changes of constraints' parameters in the optimization problem, depicting the fact that the number of robots needed to perform a given task changes during its execution.

The assignment problem is solved using any other method used in optimization problem. Here we suggest the following, however the user can replace the solver without changing the BT structure:

$$\begin{aligned} & \text{minimize} && - \sum_{i \in \mathcal{R}} \sum_{l_j \in \bar{\mathcal{L}}} p(i, l_j) r(i, l_j) \\ & \text{subject to:} && a_j \leq \sum_{i \in \mathcal{R}} r(i, l_j) \leq b_j \quad \forall l_j \in \mathcal{L}_i \\ & && \sum_{l_j \in \bar{\mathcal{L}}} r(i, l_j) \leq 1 \quad \forall i \in \mathcal{R} \\ & && r(i, j) \in \{0, 1\} \end{aligned} \quad (2)$$

where $r : \mathcal{R} \times \bar{\mathcal{L}} \rightarrow \{0, 1\}$ is a function that represents the assignment of robot i to a local task l_j , taking value 1 if the assignment is done and 0 otherwise; $p : \mathcal{R} \times \bar{\mathcal{L}} \rightarrow \bar{\mathbb{R}}$ is a function assigning a performance to robot i at executing the task l_j (if the robot i cannot perform the task l_j i.e., $l_j \notin \mathcal{L}_i$ then $p(i, l_j) = -\infty$). $a_j, b_j \in \mathbb{N}$ are respectively the minimum and the maximum number of robots assignable to the task l_j and they change during the execution of (1) and $a_j \in \mathbb{N}$ is set to a positive value upon requests from \mathcal{T}_G . At the beginning they are initialized to 0, since no assignment is needed. When the execution of a global task g_k starts, a_j and b_j are set respectively to $\nu(g_k, l_j)$ and $\mu(g_k, l_j) \quad \forall l_j \in \psi(g_k)$. Then when a local task l_j finishes its execution (i.e., it returns either success or failure), a_j and b_j decrease by 1 until $a_j = 0$, since the robot executing l_j can be assigned to another task while the other robots are still executing their local task. When the task g_k finishes, both a_j and b_j are set back to zero.

The tree \mathcal{T}_A is shown in Fig. 6(a). The condition “Local Task Finished” returns success if a robot has succeeded or failed a local task, failure otherwise. For ease of implementation a robot can communicate to all the other when has finished a local task. The condition “New Global Task Executed” returns success if it is satisfied, failure otherwise.

The condition “Check Consistency” checks if the constraints of the optimization problem are consistent with each other, since such constraints change their parameters during the tree execution. The addition of a constraint in (2) effects the system only if a solution exists. This condition is crucial to run a number of global tasks in parallel according to the number of available robots (see Example 1 below).

The action “Assign Agents” is responsible for the task assignment, it returns running while the assignment routine is executing, it returns success when the assignment problem is solved and it returns failure if the optimal value is ∞ .

b) *Definition of \mathcal{T}_G* : This tree is designed to achieve the global goal defined in Problem 1 (i.e., the global goal is achieved if the tree returns success) executing a set of global

tasks.

An example of \mathcal{T}_G is depicted in Fig. 6(b). The execution

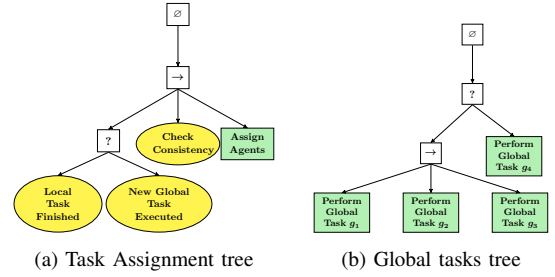


Fig. 6. Example of BT modeling a global task

of the action “Perform Global Task g_k ” requires that some robots have to be assigned, then it sets $b_j = \mu(g_k, l_j) \quad \forall l_j \in \psi(g_k)$. In the tree \mathcal{T}_A the condition “New Global Task Executed” is now satisfied, making a new assignment if the constraints are consistent with each other. Finally action “Perform Global Task g_k ” returns success if the number of local tasks $l_j \in \psi(g_k)$ that return success is greater than $\nu(g_k, l_j)$ (i.e the minimum amount of robots needed have succeeded) it returns failure if the number of local tasks in $\psi(g_k)$ that return failure is greater than $\sum_{i \in \mathcal{R}} r(i, j) - \nu(g_k, l_j)$ (i.e., some robots have failed to perform the task, the remaining ones are not enough to succeed), it returns running otherwise. The case of translating a single robot BT to a multi robot BT is described in Section V-A below.

c) *Definition of \mathcal{T}_{Li}* : This tree comprises the planner of a single agent. Basically it executes one of its subtrees upon request from the \mathcal{T}_A , the request is made by executing the action “Assign Agents”.




Fig. 7. Example of BT modeling a local task on a single agent

Example 1: By way of example, we consider a system that is to explore 2 different areas. Intuitively, different areas can be explored by different robots at the same time but if only a single robot is available those areas can only be explored in a sequence. The local trees \mathcal{T}_{Li} are as depicted in Fig. 8(b) and the global task tree \mathcal{T}_G as in Fig. 8(a).

Considering first the case in which there is only one robot available. When the execution of *Explore Area A* in \mathcal{T}_G starts, in (2) the constraint that one robot has to be assigned to explore area A is added (i.e. the related a_j and b_j are both set to 1). The optimization problem is feasible hence the assignment takes place. When the execution of *Explore Area B* starts, the related constraint is not consistent with the other constraints, since the only available robot cannot be assigned

to two tasks, hence the assignment does not take place. Considering now the case where a new robot is introduced in the system, it is possible to assign two robots. Both constraints: the one related to *Explore Area A* and the one related to *Explore Area B*, can be introduced in the optimization problem without jeopardizing its consistency. Those two tasks can be executed in parallel. Note that the two different executions (i.e. *Explore Area A* first, then *Explore Area B*; *Explore Area A* and *Explore Area B* in parallel) depend only on the number of available robots, the designed tree does not change.

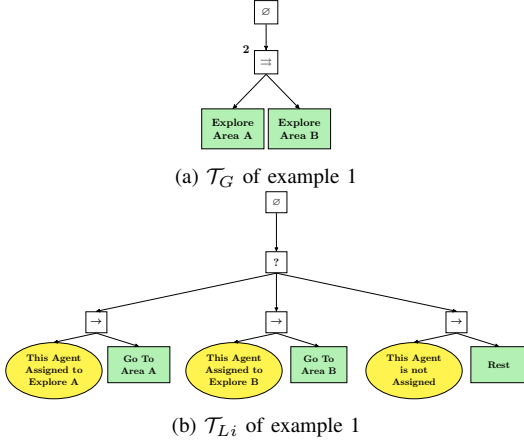


Fig. 8. Trees of example 1

A. From single robot BT to multi-robots BT

Here we present how one can design a BT for a multi-robot system reusing the BT designed to control a single robot. Let \mathcal{T}_S be the tree designed for the single robot execution with $\mathcal{A}_S = \{\alpha_1, \alpha_2, \dots\}$ being the set of action nodes, each robot in the multi-robot system will have as \mathcal{T}_G a tree with the same structure of \mathcal{T}_S changing the meaning of the action nodes and replacing the control flow nodes with parallel nodes where possible.

The action nodes of \mathcal{T}_G are the global tasks \mathcal{G} of the tree (i.e. tasks that the entire system has to perform). Their execution has the same effect on the multi-robot system as the effect of the actions \mathcal{A}_S in \mathcal{T}_S . The execution of tasks in \mathcal{G} adds constraints in the optimization problem of the tree \mathcal{T}_A defined by the sets $\psi(g_k), \forall g_k \in \mathcal{G}$. Those sets, in this case, have all cardinality 1 (i.e., only one robot is assigned to perform a task in \mathcal{G}) since it is a translation from a single robot execution to a multi-robot execution.

The main advantage of having a team of robots lies in fault tolerance and the possibility of parallel execution of some tasks. In the tree \mathcal{T}_G the user can replace sequence and selector nodes with parallel nodes and $M = N$ and $M = 1$ respectively (i.e., if the children of a selector node can be executed in parallel, the correspondent parallel node returns success as soon as one child returns success; on the other hand if the N children of a sequence node can be executed in parallel, the correspondent parallel node returns success only when all of them return success). The trees \mathcal{T}_{L_i} have

the same structure of the one defined in section V, defining as “Perform Local Task k ” the action α_k .

Then to construct \mathcal{T}_G , first each action node of \mathcal{T}_S is replaced by a global task g_k with $\psi(g_k) = \alpha_k$ and then the control flow nodes are replaced with parallel nodes where possible. The tree \mathcal{T}_A is always the one defined in section V. Finally, each robot i in a multi-robot system runs the tree \mathcal{T}_i defined in (1).

Here we observe that if the multi-robot system is heterogeneous the trees \mathcal{T}_{L_i} are such that they hold the condition $\bar{\mathcal{L}} \subseteq \mathcal{A}_S$ (i.e. the multi-robot system can perform the tasks performed by the original single robot).

B. Fault Tolerance

The fault tolerance capabilities of a multi-robot system are due to the redundancy of tasks executable by different robots. In particular, a local task l_j can be executed by every robot $i \in \mathcal{R}$ satisfying the condition $l_j \in \mathcal{L}_i$. Here we make the distinction between *minor* faults and *major* faults of a robot according to their severity. A minor fault involves a single local task l_j , i.e. the robot is no longer capable to perform the task l_j , here the *level* of a minor fault is the number of local tasks involved. A major fault of a robot implies that the robot can no longer perform any of its local task.

Remark 2: A minor fault of level λ is different from λ minor faults, since the latter might involve different robots.

Remark 3: A major fault involving a robot $i \in \mathcal{R}$ is equivalent to a minor fault of level $|\mathcal{L}_i|$ involving every task in \mathcal{L}_i .

Definition 1: A multi-robot system is said to be *weakly fault tolerant* if it can tolerate any minor fault.

Definition 2: A multi-robot system is said to be *strongly fault tolerant* if it can tolerate any major fault.

Lemma 1: A multi-robot system is weakly fault tolerant if and only if for each robot $i \in \mathcal{R}$ and for each local task $l_j \in \mathcal{L}_i$ there exists another robot $h \in \mathcal{R}$ such that $l_j \in \mathcal{L}_i \cap \mathcal{L}_h$ and the problem (2) is consistent under the constraint $r(i, j) \cdot r(h, j) = 0$.

Proof: (if) When $r(i, j) \cdot r(h, j) = 0$ either robot i or h is not deployed. Assume that the robot i is involved in a minor fault, related to the local tasks l_j . Since $l_j \in \mathcal{L}_i \cap \mathcal{L}_h$, $l_j \in \mathcal{L}_h$ then the robot h can perform the local task l_j in place of robot i , since it is not deployed.

(only if) if the system can tolerate a minor fault then for each robot i there exists another robot h such that the local task related to the fault can be performed by both of them since one took over the other. Hence $l_j \in \mathcal{L}_i$ and $l_j \in \mathcal{L}_h$, this implies that $l_j \in \mathcal{L}_i \cap \mathcal{L}_h$ holds. Moreover, if the fault can be tolerated it means that when l_j is performed either the robot i or robot h is not deployed otherwise the reassignment would not be possible. Hence the condition $r(i, j) \cdot r(h, j) = 0$ holds. ■

Corollary 1: A multi-robot system can tolerate a minor fault of level λ if Lemma 1 holds for the local tasks involved in the fault.

Lemma 2: A multi-robot system is strongly fault tolerant if and only if for each robot i there exists a set of robots I

such that $\mathcal{L}_i \subseteq \cup_{h \in I, h \neq i} \mathcal{L}_h$ and the problem (2) is consistent under the constraints $r(i, j) \cdot r(h, j) = 0 \forall j: l_j \in \mathcal{L}_i, h \in I$.

Proof: (if) Since $\mathcal{L}_i \subseteq \cup_{h \in I, h \neq i} \mathcal{L}_h$ all the local task of the robot i can be performed by some other robots in the system. Then when the robot i is asked to perform a local task l_j , exists at least another robot that can perform such task, this robot is not deployed since $r(i, j) \cdot r(h, j) = 0 \forall j: l_j \in \mathcal{L}_i, h \in I$.

(only if) if the system can tolerate a major fault then exists a robot i that is redundant, i.e. its tasks can be performed by other robots. Hence there exists a set of robots I , in which i is not contained, such that $\mathcal{L}_i \subseteq \cup_{h \in I, h \neq i} \mathcal{L}_h$, moreover since each local task l_j of the robot i can be performed by a robot $h \in I$, robot h is not deployed when i is performing l_j otherwise it could not take over robot i , hence $r(i, j) \cdot r(h, j) = 0 \forall j: l_j \in \mathcal{L}_i, h \in I$. ■

VI. MOTIVATING EXAMPLE

In this section we expand upon the example that was briefly mentioned in Section I, and involves a multi-robot system that is to replace damaged parts of the electronic hardware of a vehicle. Let us assume that to repair the vehicle, the system has to open the metal cover first, then it has to check which parts have been damaged, replace the damaged parts, solder the connecting wires, and finally close the metal cover. In this example the setup is given by a team composed by robots of type A , B and C . Robots of type A are small mobile dual arm robots that can carry spare hardware pieces, place them and remove the damaged part from the vehicles, and do diagnosis. Robots of type B are also small mobile dual arm robots with a gripper and soldering iron as end effectors. Robots of type C are big dual arm robots that can remove and fasten the vehicle cover. We consider the general case in which the vehicle's electronic hardware consists of p parts, all of them replaceable. The global tasks set is defined as:

$$\mathcal{G} = \{\text{Normal Operating Condition, Open Metal Cover, Close Metal Cover, Fix Part 1, Fix Part 2, \dots, Fix Part } p, \text{ Diagnose Part 1, Diagnose Part 2, \dots, Diagnose Part } p\}. \quad (3)$$

where:

$$\begin{aligned} \text{Open Metal Cover} &= \{\text{Remove Screws, Remove Cover}\} \\ \text{Close Metal Cover} &= \{\text{Place Screws, Place Cover}\} \\ \text{Fix Part } k &= \{\text{Fix HW } k, \text{Fix Wires } k, \\ &\quad \text{Solder } k\} \quad k = \{1, 2, \dots, p\}. \end{aligned} \quad (4)$$

The local tasks for robots of type A set are:

$$\mathcal{L}_i = \{\text{Replace HW, Replace Wires, Do Diagnosis on } k\} \quad i = \{1, 2\}. \quad (5)$$

The local tasks for robots of type B set are:

$$\mathcal{L}_i = \{\text{Replace HW, Replace Wires, Do Diagnosis on } k, \text{Solder}\} \quad i = \{3, 4\}. \quad (6)$$

Task	Robots A	Robots B	Robots C
Use Screwdriver	$-\infty$	$-\infty$	1
Move Frame	$-\infty$	$-\infty$	1
Do Diagnosis on k	3	1	$-\infty$
Replace HW k	1.5	1	$-\infty$
Replace Wires k	1.5	1	$-\infty$
Use Soldering Iron on k	$-\infty$	1	$-\infty$

TABLE I. Robot performances values

Patameter	Value
$\nu(\text{Remove Screws, Use Screwdriver})$	1
$\mu(\text{Remove Screws, Use Screwdriver})$	2
$\nu(\text{Place Screws, Use Screwdriver})$	1
$\mu(\text{Place Screws, Use Screwdriver})$	2
$\nu(\text{Remove Cover, Move Frame})$	2
$\mu(\text{Remove Cover, Move Frame})$	2
$\nu(\text{Place Screws, Remove Cover})$	1
$\mu(\text{Place Screws, Remove Cover})$	2
$\nu(\text{Place Cover, Move Frame})$	2
$\mu(\text{Place Cover, Move Frame})$	2
$\nu(\text{Diagnose Part } k, \text{ Do Diagnosis on } k)$	1
$\mu(\text{Diagnose Part } k, \text{ Do Diagnosis on } k)$	1
$\nu(\text{Fix HW } k, \text{ Replace HW } k)$	1
$\mu(\text{Fix HW } k, \text{ Replace HW } k)$	1
$\nu(\text{Fix Wires } k, \text{ Replace Wires } k)$	1
$\mu(\text{Fix Wires } k, \text{ Replace Wires } k)$	2
$\nu(\text{Solder } k, \text{ Use Soldering Iron on } k)$	1
$\mu(\text{Solder } k, \text{ Use Soldering Iron on } k)$	3

TABLE II. Scenario parameters. Note that when $\nu = \mu$, as in the case of (Remove Cover, Move Frame) we have that the number of required robots for a task is equal to the maximal number of robots that can participate in executing the task.

The global tasks for robots of type C set are:

$$\mathcal{L}_i = \{\text{Use Screwdriver, Move Frame } i = \{5, 6\}\}. \quad (7)$$

The map from global tasks to local tasks ψ are as follows:

$$\begin{aligned} \psi(\text{Remove Screws}) &= \text{Use Screwdriver} \\ \psi(\text{Place Screws}) &= \text{Use Screwdriver} \\ \psi(\text{Remove Cover}) &= \text{Move frame} \\ \psi(\text{Place Cover}) &= \text{Move frame} \\ \psi(\text{Diagnose Part } k) &= \text{Do Diagnosis on } k \\ \psi(\text{Fix HW } k) &= \text{Replace HW } k \\ \psi(\text{Fix Wires } k) &= \text{Replace Wires } k \\ \psi(\text{Solder } k) &= \text{Use Soldering Iron on } k \end{aligned} \quad (8)$$

The tree \mathcal{T}_G defined is showed in Fig. 9.

A. Execution

Here we describe the execution of the above mission using the proposed framework.

Consider a team of 2 robots of type A , 2 robots of type B , and 2 robots of type C designate to diagnose and fix a vehicle composed by 5 critical parts as depicted in Fig. 10(a). The robots' performances $p(i, l_j)$ are related to how fast tasks are accomplished. They are collected in Table I, and the scenario's parameters are in Table II.

When the vehicle is running on nominal operating conditions, the assignment tree \mathcal{T}_A of each robot has $a_j =$

$b_j = 0$ since no assignments are needed. When a fault on the vehicle takes place, the condition “Nominal Operating Conditions” is no longer true in all the robots’ BTs, then the “Remove Screws” action has to be executed. According to $\psi(\text{Remove Screws})$ this global task requires the local task “Use Screwdriver” to be accomplished. The constraints related to this task have to be added in (2) after checking the feasibility of the optimization problem. The parameter a_j and b_j for “Use Screwdriver” are changed into: $a_j = \nu(\text{Remove Screws, Use Screwdriver}) = 1$ and $b_j = \mu(\text{Remove Screws, Use Screwdriver}) = 2$ with $j : l_j = \{\text{Use Screwdriver}\}$. Solving (2) both robots C are assigned to “Use Screwdriver”, while the other are not deployed (Fig. 10(b)). When a robot accomplishes its local task, then the corresponding b_j is decreased by 1 until $b_j = 0$ and the robot can be deployed for a new task if needed. A similar assignment is then made to accomplish the global task “Remove Cover”. After removing the metal cover, all the vehicle hardware have to be diagnosed, and the global task “Diagnose Part 1” will change the parameters in $a_j = \nu(\text{Diagnose Part 1, Do Diagnosis on 1}) = 1$ and $b_j = \mu(\text{Diagnose Part 1, Do Diagnosis on 1}) = 1$ with $j : l_j = \{\text{Do Diagnosis on 1}\}$, in a similar way “Diagnose Part 2” will change the parameters in $a_j = \nu(\text{Diagnose Part 2, Do Diagnosis}) = 1$ and $b_j = \mu(\text{Diagnose Part 2, Do Diagnosis}) = 1$ with $j : l_j = \{\text{Do Diagnosis}\}$ and so forth until “Diagnose Part 4”. When the system changes the parameters related to “Diagnose Part 5” the optimization problem is not feasible since there are not enough agents, then the assignment for “Diagnose Part 5” does not take place (Fig. 10(c)).

Here we show the reactive functionality of the proposed framework. Note that the robots of type A have higher performance than the robots of type B in doing diagnosis according to Table I, i.e., they are faster. While the robots of type B are still doing diagnosis on $HW3$ and $HW4$, the robots of type A can be deployed for new tasks without waiting for the slower robots (Fig. 10(d)). After diagnosis, $HW2$ and $HW4$ must be replaced, robots of type A have higher performance in “Replace HW” therefore, they are deployed to minimize the objective function of (2), robots of type B are not deployed since $\mu(\text{Fix HW } k, \text{ Replace HW } k) = 1$ i.e. the replacing of an electronic HW is made by a single robot (Fig. 10(e)).

If a robot of type A has a major fault, all its performance parameters are set to $-\infty$ and a robot of type B can take over since it can execute all the task of a robot A , this situation is depicted in Fig. 10(f). Here two robots of type B are replacing wires on $HW1$, since $\mu(\text{Fix Part } i, \text{ Replace Wires}) = 2$ and a robot of type A is replacing wires on $HW2$. The final part of the repair is the soldering of the new wires, executable only by robots of type B (Fig. 10(g)). Finally the metal cover is put back (Fig. 10(h)) and the tree \mathcal{T}_G returns success.

B. Fault Tolerance Analysis

Analyzing the system, it is neither weakly nor strongly fault tolerant. A minor fault on the local task “Move Frame”

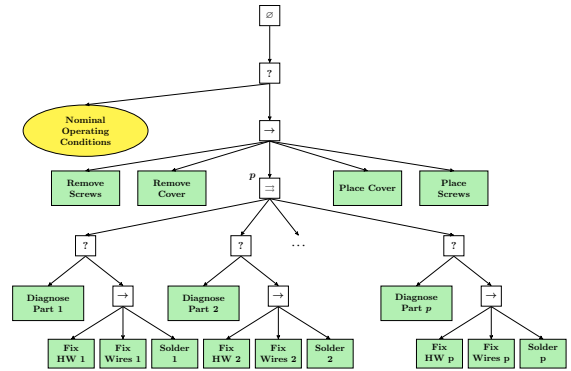


Fig. 9. \mathcal{T}_G of Example VI-A

can not be tolerated since $\nu(\text{Remove Cover, Move Frame}) = 2$ as well as $\nu(\text{Place Cover, Move Frame}) = 2$ thus a single robot cannot perform the task. However, some faults can be tolerated. The multi-robot system can tolerate up to 3 major faults (2 faults involving both robots A , and 1 fault involving a robot B) and up to 11 minor faults (1 fault involving “Use Screwdriver”, 3 faults involving “Replace HW k ”, 3 faults involving “Replace Wires k ”, and 3 faults involving “Do diagnosis on k ”), as long as at least one robot of type B and 2 robots of type C are operating. Note that the local task “Use soldering iron on k ” can be performed by a robot of type B only.

VII. CONCLUSIONS

We have shown a possible use of BTs as a distributed controller for multi-robot systems working towards global goals. This extends the fields where BTs can be used, with a strong application in robotic systems. We have also shown how using BTs as a framework for robot control becomes natural going from a plan meant to be executed by a single robot to a multi-robot plan.

ACKNOWLEDGMENT

This work has been supported by the Swedish Research Council (VR) and the European Union Project RECONFIG, (FP7-ICT-2011-9), the authors gratefully acknowledge the support.

REFERENCES

- [1] D. Isla, “Handling Complexity in the Halo 2 AI,” in *Game Developers Conference*, 2005.
- [2] A. Champandard, “Understanding Behavior Trees,” *AiGameDev.com*, vol. 6, 2007.
- [3] D. Isla, “Halo 3-building a Better Battle,” in *Game Developers Conference*, 2008.
- [4] C. Lim, R. Baumgarten, and S. Colton, “Evolving Behaviour Trees for the Commercial Game DEFCON,” *Applications of Evolutionary Computation*, pp. 100–110, 2010.
- [5] D. Perez, M. Nicolau, M. O’Neill, and A. Brabazon, “Evolving Behaviour Trees for the Mario AI Competition Using Grammatical Evolution,” *Applications of Evolutionary Computation*, 2011.
- [6] A. Shoulson, F. M. Garcia, M. Jones, R. Mead, and N. I. Badler, “Parameterizing Behavior Trees,” in *Motion in Games*. Springer, 2011.
- [7] I. Bojic, T. Lipic, M. Kusek, and G. Jezic, “Extending the JADE Agent Behaviour Model with JBehaviourtrees Framework,” in *Agent and Multi-Agent Systems: Technologies and Applications*. Springer, 2011, pp. 159–168.

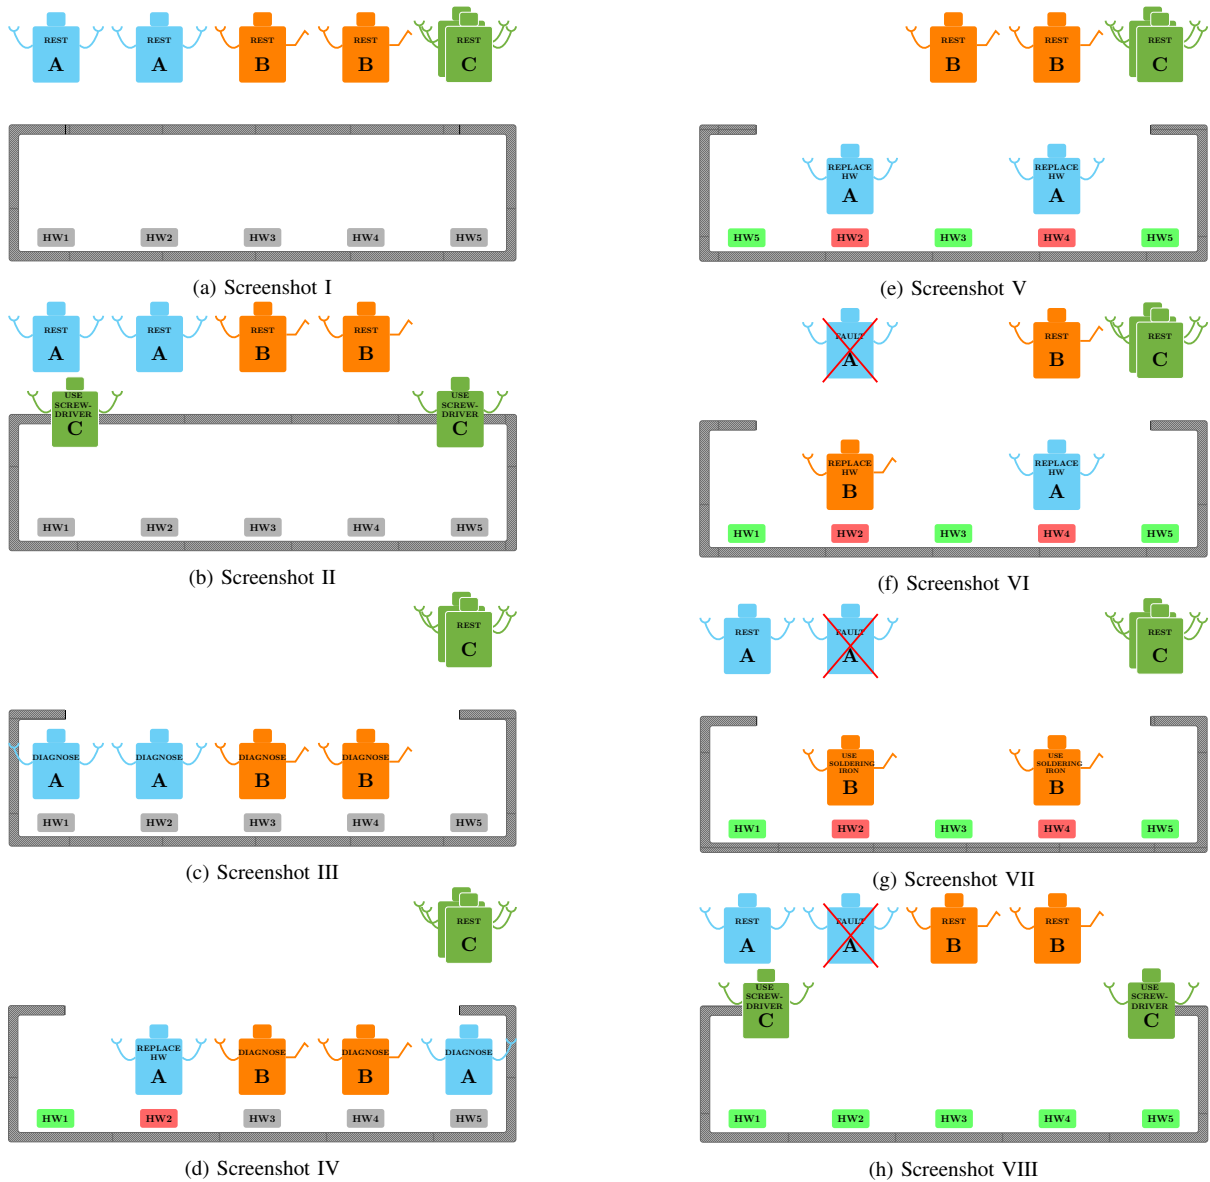


Fig. 10. Plan Execution. HW components are either un-diagnosed (grey), checked and OK (green) or broken (red)

- [8] P. Ögren, "Increasing Modularity of UAV Control Systems using Computer Game Behavior Trees," in *AIAA Guidance, Navigation and Control Conference*, Minneapolis, MN, 2012.
- [9] J. A. D. Bagnell, F. Cavalcanti, L. Cui, T. Galluzzo, M. Hebert, M. Kazemi, M. Klingensmith, J. Libby, T. Y. Liu, N. Pollard, M. Pivtoraiko, J.-S. Valois, and R. Zhu, "An Integrated System for Autonomous Robotics Manipulation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2012, pp. 2955–2962.
- [10] A. Klöckner, "Interfacing Behavior Trees with the World Using Description Logic," in *AIAA conference on Guidance, Navigation and Control*, Boston, 2013.
- [11] M. Colledanchise and P. Ögren, "How Behavior Trees Modularize Robustness and Safety in Hybrid Systems," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, Sept 2014, pp. 1482–1488.
- [12] E. W. Dijkstra, "Letters to the editor: go to statement considered harmful," *Commun. ACM*, vol. 11, pp. 147–148, March 1968.
- [13] X. C. Ding, M. Kloetzer, Y. Chen, and C. Belta, "Automatic deployment of robotic teams," *Robotics Automation Magazine, IEEE*, vol. 18, no. 3, pp. 75–86, 2011.
- [14] S. Karaman and E. Frazzoli, "Vehicle routing problem with metric temporal logic specifications," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, 2008, pp. 3953–3958.
- [15] A. Ulusoy, S. L. Smith, and C. Belta, "Optimal multi-robot path planning with ltl constraints: Guaranteeing correctness through synchronization," *CoRR*, vol. abs/1207.2415, 2012.
- [16] C. Castelfranchi and R. Falcone, "Towards a Theory of Delegation for Agent-based Systems," *Robotics and Autonomous Systems*, vol. 24, pp. 141–157, 1998.
- [17] R. Falcone and C. Castelfranchi, "C.: The human in the loop of a delegated agent: The theory of adjustable social autonomy," *IEEE Transactions on Systems, Man and Cybernetics Part A: Systems and Humans*, pp. 406–418, 2001.
- [18] P. Doherty, D. Landn, and F. Heintz, "A distributed task specification language for mixed-initiative delegation," 2012.
- [19] A. Marzinotto, M. Colledanchise, C. Smith, and P. Ögren, "Towards a Unified Behavior Trees Framework for Robot Control," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, June 2014.
- [20] M. Colledanchise, A. Marzinotto, and P. Ögren, "Performance Analysis of Stochastic Behavior Trees," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, June 2014.