Improving the Parallel Execution of Behavior Trees

Michele Colledanchise and Lorenzo Natale

Abstract— Behavior Trees (BTs) have become a popular framework for designing controllers of autonomous agents in the computer game and in the robotics industry. One of the key advantages of BTs lies in their modularity, where independent modules can be composed to create more complex ones. In the classical formulation of BTs, modules can be composed using one of the three operators: Sequence, Fallback, and Parallel. The Parallel operator is rarely used despite its strong potential against other control architectures such as Finite State Machines. This is due to the fact that concurrent actions may lead to unexpected problems similar to the ones experienced in concurrent programming. In this paper, we outline how to tackle the aforementioned problem by introducing Concurrent BTs (CBTs) as a generalization of BTs in which we include the notions of progress and resource usage. We show how CBTs allow safe concurrent executions of actions and we analyze the approach from a mathematical standpoint. To illustrate the use of CBTs, we provide a set of use cases in realistic robotics scenarios.

I. INTRODUCTION

Behavior Trees (BTs) were introduced in the computer game industry to model the behavior of Non-Player Characters (NPCs) [1]–[3]. BTs have now matured to the level that they are included in several textbooks on the topic [4], [5] and game engines. Following the development in computer game industry, BTs now also finds large use in real-robot applications, including manipulation [6], non-expert programming [7], brain surgery [8]. Other works include UAV missions [9], [10], the computer verification of plans [9], and the estimation of resulting execution times in [11].

BTs are appreciated because they are modular, flexible, and reusable, and have also been shown to generalize successful control architectures such as the Subsumption architecture [12] and the Teleo-reactive Paradigm [13] in [14].

The Parallel operator in BTs has the great advantage that it helps to tame the *curse of dimensionality* that affects many other control architectures, in which the system's complexity is the product of its sub-systems' complexities [15]. However, in the classical formulation of BTs, the trees in a Parallel operator are independent from each other. This limits their applications to orthogonal tasks [14], [16], making the Parallel operator rarely used compared to the other ones.

We propose Concurrent BTs (CBTs) as a generalization of BTs where for each node we define its progress value and the resources needed for its execution. We show how CBTs allow actions to be safely executed concurrently by adopting

(a) Example of two BTs executed in parallel that require process synchronization.

(b) Example of two BTs executed in parallel that require resource synchronization.

Fig. 1. Example of two BTs executing tasks in parallel that require synchronization.

solutions from computer programming in terms of process synchronization and exclusive use of critical resources, as in the two following examples.

1) A robot that is asked to find and recognize the objects on the floor. A possible way to describe this task is using the BT in Figure 1(a). Where the action *Navigate* generates and makes the robot follow a path and the action *Seek* looks for objects on the floor and recognizes them. However, since classical formulation of BTs executes the two actions independently, there could be the case where the robot navigates in the room too fast for the seeking routine to recognize objects, resulting in an unsuccessful execution of the task.

2) A service robot, beside its main task, has also to monitor the light condition of the room and whether a person gets too close to the robot. A possible way to describe this monitoring task is using the BT in Figure 1(b). Each task has a warning message that invokes the speaker. However, the speaker allows one speech at a time. For this minor action, the tasks cannot be considered orthogonal and cannot be safely executed in parallel using the classical formulation of BTs.

It is clear that the parallel execution of BTs needs synchronization strategies for exploiting its full potential. We believe that proper parallel task executions will bring benefits in terms of efficiency and multitasking to robot programming in a similar way as done in computer programming. However, in allowing arbitrary BTs to be executed in parallel we inherit fundamental problems of parallel computing: *process synchronization*, where a process has to wait for another process for the execution, and *data synchronization*, where the access to critical resources has to be regulated.

It is important to stress that execution of some routines may require different amount of time depending on the context (e.g. object recognition may be triggered only when visual inspection spots a potential object). Moreover there could be a substantial difference between the execution of an action in a simulation and in the real world. Therefore

synchronization must be ensured at run-time.

The contributions of this paper are the definition of CBTs, in the Parallel composition of CBTs for process and resource synchronization, and a formal mathematical analysis. To the best of our knowledge, this is the first attempt at synchronizing the concurrent executions of BTs as it is done in computer programming.

The remainder of this paper is structured as follows: In Section II we overview the related work. In Section III we present the classical formulation of BTs. Then, in Section IV, we formulate the problem and in Section V we show the proposed solution. In Section VII some practical observations and in Section VI we analyze the solution from a mathematical standpoint. Finally, in Section VIII, we present some use cases. We conclude the paper in Section IX.

II. RELATED WORK

The Parallel operator is the least used operator in BT applications as it entails concurrency problems (e.g. race conditions, starvation, deadlocks, etc). Their applications usually assume one of the following: 1) the actions executed in parallel lie on orthogonal state spaces [16] or 2) the actions that can be potentially executed in parallel have a predefined priority assigned [17].

In [18] the authors use a Parallel operator to monitor user requests and execute an activity accordingly. Since the user can define one activity at a time, the activities are mutually exclusive. Hence the BT for each activity (and its actions) is defined in its own sub-space.

In [15] the authors show how to improve fault tolerance and performance of a single robot BT by adding more robots and extending the BT into a so-called *multirobot* BT. The authors show that, under some assumptions, by turning Sequence and Fallback operators of the original BT into a Parallel we can improve task performance. The Parallel operator involves multiple agents, each assigned to a specific task by a task-assignment algorithm. The task-assignment algorithm ensures the absence of conflicts.

A more recent work [19] uses BTs to represent medical procedures. The Parallel operator is used to monitor concurrent procedures that are not in conflict by definition. Hence each procedure operates in its own sub-space

A BT-based task planner that makes large use of the Parallel operator is the A Behavior Language (ABL) [20]. ABL was originally designed for the dialogue game *Façade*, and it is appreciated for its ability to handle planning and acting on multiple scales as, in particular, Real-Time Strategy games [17]. ABL executes sub-plans in parallel and resolves conflicts between multiple concurrent actions by defining a fixed priority order. This jeopardizes the modularity of BTs and introduces the threat of starvation (i.e. the execution of an action with low priority is perpetually denied).

To conclude, there is currently no work addressing the synchronization problems of the Parallel operator. This makes our paper fundamentally different than the ones presented above and the BT literature.

III. BACKGROUND:BEHAVIOR TREES AND CONCURRENT PROGRAMMING

In this section, we briefly describe the semantic of classical BTs and the standard semantics for concurrent programming. A more detailed description of BTs can be found in [14] while a more detailed description of concurrent programming can be found in [21].

A. Behavior Trees

A BT is a graphical modeling language used as a representation for execution of observation-based actions in a system. A BT is represented as a rooted tree where the internal nodes represent operators and leaf nodes represent actuation or sensing skills.

Graphically, the children of a node are placed below it and they are executed in the order from left to right, as will be explained later.

The execution of a BT begins from the root node. It sends activation signals called *Ticks* with a given frequency to its children. A node in the tree is executed if and only if it receives Ticks from its parent. When the node no longer receives Ticks, its execution is aborted. The child returns to the parent a status *Success*, *Running*, or *Failure* according to the node's logic. Below we present the most common BT nodes.

When a Fallback operator receives Ticks, it routes them to its own children from the left, returning Success/Running as soon as it finds a child that returns Success/Running. It returns Failure when all the children return Failure. When a child returns Running or Success, the Fallback operator does not send Ticks the next child (if any). The Fallback operator is graphically represented by a box with the label "?", as in Figure 2, and its pseudocode is described in Algorithm 1

When a Sequence operator receives Ticks, it routes them to its own children from the left, returning Failure/Running as soon as it finds a child that returns Failure/Running. It returns Success when all the children return Success. When a child returns Running or Failure, the Sequence node does not send Ticks the next child (if any).The Sequence operator is graphically represented by a box with the label "→", as in Figure 2, and its pseudocode is described in Algorithm 2

The Parallel operator ticks its children in parallel and returns Success if all children return Success, it returns Failure if at least one child returns Failure, and it returns Running otherwise. The Parallel operator is graphically represented by a box with the label "⇉", as in Figure 2, and its pseudocode is described in Algorithm 3.

Remark 1: The for loop in Algorithm 3 Line 2 executes the loops in parallel.

As long as an Action receives Ticks, it performs some operations. It returns Success to its parent if the operations are completed and Failure if the operations cannot be completed. Otherwise, it returns Running. Whenever a running Action does no longer receive Ticks, its execution is aborted. An Action is graphically represented by a rectangle, as in Figure 2, and its pseudocode is described in Algorithm 4.

Algorithm 1: Pseudocode of a Fallback operator with N children

1 Function *Tick()*
2   **for** $i \leftarrow 1$ **to** N **do**
3   *childStatus* $\leftarrow$ *child(i)*.Tick()
4   **if** *childStatus = Running* **then**
5   **return** *Running*
6   **else if** *childStatus = Success* **then**
7   **return** *Success*
8   **return** *Failure*

Algorithm 2: Pseudocode of a Sequence operator with N children

1 Function *Tick()*
2   **for** $i \leftarrow 1$ **to** N **do**
3   *childStatus* $\leftarrow$ *child(i)*.Tick()
4   **if** *childStatus = Running* **then**
5   **return** *Running*
6   **else if** *childStatus = Failure* **then**
7   **return** *Failure*
8   **return** *Success*

Algorithm 3: Pseudocode of a Parallel operator with N children

1 Function *Tick()*
2   **forall** $i \leftarrow 1$ **to** N **do**
3   *childStatus*[i] $\leftarrow$ *child(i)*.Tick()
4   **if** $\Sigma_{i:childStatus[i]=Success}1 = N$ **then**
5   **return** *Success*
6   **else if** $\Sigma_{i:childStatus[i]=Failure}1 > 0$ **then**
7   **return** *Failure*
8   **else**
9   **return** *Running*

Algorithm 4: Pseudocode of a BT Action

1 Function *Tick()*
2   *DoAStepOfComputation()*
3   **if** *action-succeeded* **then**
4   **return** *Success*
5   **else if** *action-failed* **then**
6   **return** *Failure*
7   **else**
8   **return** *Running*

Algorithm 5: Pseudocode of a BT Condition

1 Function *Tick()*
2   **if** *condition-true* **then**
3   **return** *Success*
4   **else**
5   **return** *Failure*

Whenever a Condition receives Ticks, it checks if a proposition is satisfied or not, returning Success or Failure accordingly. A Condition is graphically represented by an ellipse, as in Figure 2, and its pseudocode is described in Algorithm 5.

Remark 2: Algorithm 4 does a step of computation at each Tick. This implementation is preferred in BT libraries for computer games as Unreal Engine and PyGame. BT libraries for robotics applications prefer an implementation that allows actions to continue as long as they succeed or fails such as YARP-BT or ROS-BT. Note that in this case whenever a running action no longer receives a tick, its execution is aborted. This requires the implementation specific routine for aborting the action safely.

The state space formulation of BTs [18] allows us to study BTs from a mathematical standpoint. In that formulation, the Tick is represented by a recursive function call that includes both the return status, the system dynamics, and the system state.

Definition 1 (Behavior Tree [18]): A BT is a three-tuple

$$\mathcal{T}_i = \{f_i, r_i, \Delta t\}, \tag{1}$$

where $i \in \mathbb{N}$ is the index of the tree, $f_i : \mathbb{R}^n \to \mathbb{R}^n$ is the right hand side of a difference equation, Δt is a time step and $r_i : \mathbb{R}^n \to \{\mathcal{R}, \mathcal{S}, \mathcal{F}\}$ is the return status that can be equal to either *Running* ($\mathcal{R}$), *Success* ($\mathcal{S}$), or *Failure* ($\mathcal{F}$).

Let the Running/Activation region (R_i), Success region (S_i) and Failure region (F_i) correspond to a partitioning of the state space, defined as follows:

$$R_i = \{x : r_i(x) = \mathcal{R}\} \tag{2}$$
$$S_i = \{x : r_i(x) = \mathcal{S}\} \tag{3}$$
$$F_i = \{x : r_i(x) = \mathcal{F}\}. \tag{4}$$

Finally, let $x_k = x(t_k)$ be the system state at time t_k, then the execution of a BT $\mathcal{T}_i$ is a standard ordinary difference equation

$$x_{k+1} = f_i(x_k), \tag{5}$$
$$t_{k+1} = t_k + \Delta t. \tag{6}$$

The return status r_i will be used when combining BTs recursively, as explained below.

Having defined a BT using a state space formulation, Sequence, Fallback, and Parallel compositions are defined as follows:

Definition 2 (Sequence compositions of BTs [18]): Two or more BTs can be composed into a more complex BT using a Sequence operator,

$$\mathcal{T}_0 = \text{Sequence}(\mathcal{T}_1, \mathcal{T}_2).$$

Then r_0, f_0 are defined as follows

$$\text{If } x_k \in S_1 \tag{7}$$
$$r_0(x_k) = r_2(x_k) \tag{8}$$
$$f_0(x_k) = f_2(x_k) \tag{9}$$
$$\text{else}$$
$$r_0(x_k) = r_1(x_k) \tag{10}$$
$$f_0(x_k) = f_1(x_k). \tag{11}$$

Definition 3 (Fallback compositions of BTs [18]): Two or more BTs can be composed into a more complex BT using a Fallback operator,

$$\mathcal{T}_0 = \text{Fallback}(\mathcal{T}_1, \mathcal{T}_2).$$

Then r_0, f_0 are defined as follows

$$\text{If } x_k \in F_1 \tag{12}$$
$$r_0(x_k) = r_2(x_k) \tag{13}$$
$$f_0(x_k) = f_2(x_k) \tag{14}$$
$$\text{else}$$
$$r_0(x_k) = r_1(x_k) \tag{15}$$
$$f_0(x_k) = f_1(x_k). \tag{16}$$

Definition 4 (Parallel compositions of BTs [18]): Two or more BTs can be composed into a more complex BT using a Parallel operator,

$$\mathcal{T}_0 = \text{Parallel}(\mathcal{T}_1, \mathcal{T}_2).$$

Let $x_k = (\tilde{x}_{1k}, \tilde{x}_{2k})$ be a partitioning of the state space such that $\tilde{x}_{1k} \perp \tilde{x}_{2k}$, then r_0, f_0 are defined as follows

$$r_0(x_k) = \begin{cases} \mathcal{S} & \text{If } r_1(\tilde{x}_{1k}) = \mathcal{S} \wedge r_2(\tilde{x}_{2k}) = \mathcal{S} \\ \mathcal{F} & \text{If } r_1(\tilde{x}_{1k}) = \mathcal{F} \vee r_2(\tilde{x}_{1k}) = \mathcal{F} \\ \mathcal{R} & \text{else} \end{cases} \tag{17}$$
$$f_0(x_k) = (f_1(\tilde{x}_{1k}), f_2(\tilde{x}_{2k})) \tag{18}$$

B. Concurrent Programming

Concurrent programming refers to the execution of multiple processes during overlapping periods of time concurrently instead of sequentially.

The need for synchronization arises in any kind of concurrent process, even in single-processor systems. The main needs for synchronization are: *producer-consumer relationship*, where a consumer process has to wait until a producer produces the necessary data, and *exclusive use of resources*, where multiple processes want to access a critical resource. A synchronization strategy must ensure that only one process at a time can access the resource [21].

Barriers and *Semaphores* are popular ways to implement synchronization. A Barrier is a method that allows concurrent processes to wait for each other at a specific point of execution. A Semaphore is a method that regulates the access to critical resources.

Deadlocks and *Starvation* are among the most common threats in concurrent programming [21]. A Deadlock occurs when processes in a group are all waiting for resources that are being held by the others. In this case, the processes just keep waiting and execute no further; There are several techniques to avoid deadlocks [21]. A Starvation occurs when a process is waiting to use shared resource, but other processes monopolize it (e.g. because they have higher priority), and the first process is forced to wait indefinitely. A solution to starvation is to implement *aging*, where all process are preemptable and a process waiting a long time to use a resource gradually increases its chances to use it.

IV. PROBLEM FORMULATION

In this section, we first make a set of assumptions and definitions, then state the main problems.

To be able to synchronize the execution of concurrent BTs, they need to indicate their progress and the resources used. Hence we make the following assumptions.

Assumption 1: For each BT node there exists a function that indicates the progress of the node's execution at each state.

Assumption 2: For each BT node there exists a function that indicates the resources used in its execution.

To guarantee that only the ticked actions are executed, whenever a running action no longer receives ticks it stops its execution within a negligible time with respect to the time step Δt of Definiton 1. Hence we make the following assumption.

Assumption 3: Each action is able to stop its execution within $\tau << \Delta t$ when it no longer receives ticks. During τ the node does not change its progress value or resources needed.

We allow disturbance to affect the progress value. However, we assume that the disturbance acts only on nodes while they are running. Hence we make the following assumption.

Assumption 4: A BT can change its progress only while it is running.

For convergence, we need to state the following assumptions.

Assumption 5: An action node will terminate its execution in finite time.

Definition 5 (Progress independent CBTs): Two CBTs, $\mathcal{T}_1$ and $\mathcal{T}_2$, are said progress independent if the progress of a BT does not depend on the progress of the other.

We are now ready to formulate the main problems:

Problem 1: Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two progress independent BTs. Execute the two BTs concurrently such that $\mathcal{T}_1$ and $\mathcal{T}_2$ are both running only when they have the same progress, otherwise the BT with the highest progress waits.

Remark 3: $\mathcal{T}_1$ and $\mathcal{T}_2$ needs to be progress independent to avoid conflicts with the progress dependence that will be introduced.

Problem 2: Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two BTs. Execute the two BTs concurrently in a deadlock-free and starvation-free fashion such that they do not use the same resources when they are running.

Remark 4: Problem 2 relaxes the assumption that two CBTs operate on orthogonal state spaces, as stated in Definition 4.

V. PROPOSED SOLUTION

In this section, we present the proposed solution to Problems 1 and 2. We first define the progress and the resource function, needed to formally define CBTs. We then propose two new parallel operators to solve Problems 1 and 2. We define the operators by theirs pseudocode and their space state formulation, following the formulation of [18].

Definition 6 (Progress Function): The function $p : R^n \rightarrow [0, 1]$ is the progress function. It indicates the progress of the BT's execution at each state.

Definition 7 (Resources): L is a collection of symbols that represents the resources available in the system.

Definition 8 (Resource Function): The function $q : R^n \rightarrow 2^L$ is the resource function. It indicates the set of resources needed for a BT's execution at each state.

We can now define a CBT as BT with information regarding its progress and the resources needed as follows:

Definition 9 (Concurrent BTs): A CBT is a tuple

$$\mathcal{T}_i = \{f_i, r_i, \Delta t, p_i, q_i\}, \quad (19)$$

where i, f_i, Δt, r_i are defined as in Definition 1, p_i is a progress function, and q is the resource function.

A CBT has the functions p_i and q_i in addition to the others of Definition 1. These functions are user-defined for Actions and Condition. For the classical operators, the functions are defined below.

Definition 10 (Sequence compositions of CBTs): Two CBTs can be composed into a more complex CBT using a Sequence operator,

$$\mathcal{T}_0 = \text{Sequence}(\mathcal{T}_1, \mathcal{T}_2).$$

The functions r_0, f_0 match those introduced in Definition 2, while the functions p_0, q_0 are defined as follows

$$\text{If } x_k \in S_1 \quad (20)$$
$$p_0(x_k) = 0.5 + \frac{p_2(x_k)}{2} \quad (21)$$
$$q_0(x_k) = q_2(x_k) \quad (22)$$
$$\text{else}$$
$$p_0(x_k) = \frac{p_1(x_k)}{2} \quad (23)$$
$$q_0(x_k) = q_1(x_k). \quad (24)$$

Definition 11 (Fallback compositions of CBTs): Two CBTs can be composed into a more complex CBT using a Fallback operator,

$$\mathcal{T}_0 = \text{Fallback}(\mathcal{T}_1, \mathcal{T}_2).$$

The functions r_0, f_0 are defined as in Definition 3, while the functions p_0, q_0 are defined as follows

$$\text{If } x_k \in F_1 \quad (25)$$
$$p_0(x_k) = p_2(x_k) \quad (26)$$
$$q_0(x_k) = q_2(x_k) \quad (27)$$
$$\text{else}$$
$$p_0(x_k) = p_1(x_k) \quad (28)$$
$$q_0(x_k) = q_1(x_k). \quad (29)$$

Definition 12 (Parallel compositions of CBTs): Two CBTs can be composed into a more complex CBT using a Parallel operator,

$$\mathcal{T}_0 = \text{Parallel}(\mathcal{T}_1, \mathcal{T}_2).$$

The functions r_0, f_0 are defined as in Definition 4, while the functions p_0 and q_0 are defined as follows

$$p_0(x_k) = \min(p_1(x_k), p_2(x_k)) \quad (30)$$
$$q_0(x_k) = q_1(x_k) \cup q_2(x_k) \quad (31)$$
$$\quad (32)$$

Remark 5: Conditions nodes do not perform any action. Hence their progress function can be defined as $p(x_k) = 1$.

A. Process Synchronization of CBTs

In this section, we show how by using CBTs we can synchronize the progress of two or more actions executed in parallel, solving Problem 1.

In the classical formulation, the Parallel operator (see Definition 4) requires that the BTs are not dependent on each other, hence it is not possible to synchronize their execution. Having extended the BTs formulation with the progress function, we can define a Synchronized Parallel operator that takes into account the BTs progress when deciding where to route the Ticks.

Definition 13 (Synchronized Parallel): Two CBTs can be composed into a more complex BT using a Synchronized Parallel operator,

$$\mathcal{T}_0 = \text{ParallelSync}(\mathcal{T}_1, \mathcal{T}_2).$$

Let $x_k = (\tilde{x}_{1k}, \tilde{x}_{2k})$ be the partitioning as in Definition 4 then f_0 is defined as follows:

$$f_0(x_k) = (\tilde{f}_1(\tilde{x}_{1k}), \tilde{f}_2(\tilde{x}_{2k})) \quad (33)$$
$$\tilde{f}_i(x_k) = \begin{cases} f_i(\tilde{x}_i) & \text{if } \alpha_i(x_k) = 1 \\ \tilde{x}_i & \text{otherwise} \end{cases} \quad (34)$$
$$\alpha_1(x_k) = \begin{cases} 0 & \text{if } p_1(\tilde{x}_{1k}) > p_2(\tilde{x}_{2k}) \\ 1 & \text{otherwise} \end{cases} \quad (35)$$
$$\alpha_2(x_k) = \begin{cases} 0 & \text{if } p_2(\tilde{x}_{2k}) > p_1(\tilde{x}_{1k}) \\ 1 & \text{otherwise} \end{cases} \quad (36)$$

Algorithm 6 presents the pseudocode of a Synchronized Parallel operator. The main difference with the classical Parallel operator (Algorithm 3) lies in the addition of Lines 2-4 and Line 6, which implement the fact that a child receives Ticks only if its progress does not exceed the minimum progress. We graphically represent this operator with box with the label "$\overset{S}{\Rightarrow}$".

Remark 6: The Synchronized Parallel operator intrinsically implements a barrier (as described in Section III-B) at the minimum progress value, where all children wait for the one that has the "slowest" value to proceed.

We are now ready the see the execution of the BT in Figure 1(a) formalized in the example below.

Example 1: An object-seeking robot has to recognize objects on the floor. The robot's behavior is described by

Algorithm 6: Pseudocode of a ParallelSync node with N children

1   **Function** *Tick()*
2     $minProgress \leftarrow 1$
3   **for** $i \leftarrow 1$ **to** N **do**
4       $minProgress \leftarrow \min(minProgress, p_i)$
5   **forall** $i \leftarrow 1$ **to** N **do**
6   **if** $p_i \leq minProgress$ **then**
7         $childStatus[i] \leftarrow child(i).\texttt{Tick()}$
8   **if** $\Sigma_{i:childStatus[i]=Success}1 = N$ **then**
9   **return** *Success*
10   **else if** $\Sigma_{i:childStatus[i]=Failure}1 > 0$ **then**
11   **return** *Failure*
12   **return** *Running*

Algorithm 7: Pseudocode of a ParallelMutex node with N children

1   $\tilde{R} \leftarrow \emptyset$
2   $maxPriority \leftarrow 0$
3   **for** $i \leftarrow 1$ **to** N **do**
4     $maxPriority \leftarrow \max(maxPriority, \pi_i)$
5   **forall** $i \leftarrow 1$ **to** N **do**
6   **if** $q_i \nsubseteq \tilde{R}$ **or** $\pi_i = maxPriority$ **then**
7       $\tilde{R} \leftarrow \tilde{R} \cup q_i$
8       $childStatus[i] \leftarrow Children(i).\texttt{Tick()}$
9     $\tilde{R} \leftarrow \emptyset$
10   **if** $\Sigma_{i:childStatus[i]=Success}1 \geq N$ **then**
11   **return** *Success*
12   **else if** $\Sigma_{i:childStatus[i]=Failure}1 > 0$ **then**
13   **return** *Failure*
14 **return** *Running*

---

the BT in Figure 1(a). The progress of both actions is 0 at the beginning of the hallway and 1 when the hallway is completely navigated or sought, it is a value in $(0,1)$ according to the percentage of the hallway navigated or swept. If the robot finds an object in the middle of the hallway where progress for both actions is 0.5. Since it takes time to compute the perception routine that allows the robot to recognize the object, the progress of *Seek* is equal to 0.5 until the robot recognizes the object. The navigation action could continue its progress. However, as soon as its progress surpasses the value 0.5 the navigation no longer receives Ticks (Algorithm 6 Lines 6-7) and the execution is aborted, making the navigation stop. Once the robot recognizes the object, the progress of *Seek* increases. As soon as the progress of *Navigate* is equal to the progress of *Seek* the navigation resumes.

B. Resource Synchronization of CBTs

In this section we show how using CBTs we can execute multiple actions in parallel without resource conflicts.

Definition 14 (Mutually Exclusive Parallel): Two BTs can be composed into a more complex BT using a Mutually Exclusive Parallel operator,

$$\mathcal{T}_0 = \text{ParallelMutex}(\mathcal{T}_1, \mathcal{T}_2).$$

f_0 and q_0 are defined as follows:

$$f_0(x_k) = f_1(x_k)\alpha_1(x_k) + f_2(x_k)\alpha_2(x_k) \quad (37)$$

$$q_0(x_k) = \bigcup_{i:\alpha_i(x_k)=1} q_i(x_k) \quad (38)$$

$$\alpha_1(x_k) = \begin{cases} 1 & \text{if } q_1(x_k) \cap q_2(x_k) = \emptyset \ \vee \\ & \vee \pi_1(x_k) \geq \pi_2(x_k) \\ 0 & \text{otherwise} \end{cases} \quad (39)$$

$$\alpha_2(x_k) = \begin{cases} 1 & \text{if } q_1(x_k) \cap q_2(x_k) = \emptyset \ \vee \\ & \vee \pi_2(x_k) > \pi_1(x_k) \\ 0 & \text{otherwise} \end{cases} \quad (40)$$

$$\quad (41)$$

where $\pi_i(x_k)$ a user-defined priority assigned by any policy that implements *aging* (see Section III-B).

Algorithm 7 presents the pseudocode of a Mutually Exlusive Parallel operator. The main difference with the classical Parallel operator (Algorithm 3) lies in Lines 1-7 and Line 9, which implement the fact that a child nodes receives ticks only if it uses a resource that is not required by any currently running node. We graphically represent the operator with a box with the label "$\stackrel{M}{\Rightarrow}$".

Remark 7: As mentioned in Section III-B, using a policy that implements aging avoids starvation.

We now ready the describe the execution of the BT in Figure 1(b), formalized in the example below.

Example 2: A delivery robot is asked to distribute mail to an office's employee. During its execution, the robot monitors the possible position of people around to avoid harming them. The robot's behavior is described by the BT in Figure1(b). If a person is obstructing any obstacle-free path, it asks the person to move away. Moreover, if the room is too dark it asks to improve the light condition, for better robot's perception. Without loss of generality, let's assume that a person gets in front of the robot before this realizes that the room is too dark. ParallelMutex has $\tilde{R} = \emptyset$ (Algorithm 7 Line 1) and it ticks the action *Monitor People* first. Hence its set $q(x_k)$ (set of resource used) contains the symbol *speaker* making $\tilde{R}$ also contain the symbol *speaker* (Algorithm 7 Line 5). When the robot realizes that the room is too dark, the action *Monitor Light* is not ticked as its set $q(x_k)$ contains the symbol *speaker* and the condition in Algorithm 7 Line 4 is not satisfied (i.e. the speaker is not available at that time).

VI. THEORETICAL ANALYSIS

In this section, we prove that the proposed approach solves Problems 1 and 2.

Proposition 1: Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two progress independent BTs, $\mathcal{T}_0 = \text{ParallelSync}(\mathcal{T}_1, \mathcal{T}_2)$ solves Problem 1.

Proof: The execution of $\mathcal{T}_0$ is defined by Algorithm 6. For each state $x_k \in R^n$, one of the following occurs:

- $p_1(x_k) = p_2(x_k)$: in this case both BTs are executed as they both receive Ticks.
- $p_1(x_k) < p_2(x_k)$: Since in this case *minProgress* (Algorithm 6 Line 4) is $p_1(x_k)$ and $\mathcal{T}_2$ no longer receive Ticks (Line 6) making the execution of $\mathcal{T}_2$ stop, hence $p_2(x_{k+1}) = p_2(x_k)$.
- $p_2(x_k) < p_1(x_k)$: $p_1(x_{k+1}) = p_1(x_k)$ holds similarly as above. ∎

Moreover, by Assumptions 3 and 4 a BT stops increasing or decreasing its progress whenever it no longer receives ticks. Hence, $\mathcal{T}_1$ and $\mathcal{T}_2$ are both running only if they have the same progress, otherwise the BT with the higher progress waits for the other.

Lemma 1: Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two progress independent BTs, the execution of $\mathcal{T}_0 = \text{ParallelMutex}(\mathcal{T}_1, \mathcal{T}_2)$ is deadlock-free.

Proof: By Assumption 5 all the actions either terminate or set the value of the resource function to the empty set. Hence the resources are released in finite time. ∎

Lemma 2: Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two progress independent BTs, the execution of $\mathcal{T}_0 = \text{ParallelMutex}(\mathcal{T}_1, \mathcal{T}_2)$ is starvation-free.

Proof: In Algorithm 7 assign resources using a policy that implements aging. ∎

Proposition 2: Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two progress independent BTs, $\mathcal{T}_0 = \text{ParallelMutex}(\mathcal{T}_1, \mathcal{T}_2)$ solves Problem 2.
Proof: The execution of $\mathcal{T}_0$ is defined by Algorithm 7. For each state $x_k \in R^n$, one of the following occur:

- $q_1(x_k) \cup q_2(x_k) = \emptyset$: in this case the condition in Algorithm 7 Line 2, trivially holds for both children, which are ticked.
- $q_1(x_k) \cup q_2(x_k) \neq \emptyset$: Only the child with highest priority is executed.

hence $\mathcal{T}_1$ and $\mathcal{T}_2$ do not use the same resources when they are running. By Lemma 1 and Lemma 2 the execution of $\mathcal{T}_0$ is deadlock-free and starvation-free. ∎

VII. IMPLEMENTATION DETAILS

To allow CBTs to be implemented, we must make some key observations.

A. Wait and Halt Routines

In the formulation above, a running action no longer receives Ticks in one of these two cases: when the parent is routing the Ticks elsewhere (i.e. it no longer requires the execution of that action) and when the action has to wait for another one to reach the progress barrier (in case of Synchronized Parallel) or some resources (in case of Mutually Exclusive Parallel). As mentioned in Section III, an action node may require to execute a specific routine whenever the nodes no longer receive ticks. It is fundamental to define two different routines, one when the action is no longer required (*halt*) and one when the action has to wait

---

Algorithm 8: Pseudocode of a ParallelSync node with N children with explicit handling of the Halt and Pause routine.

```
1  Function Tick()
2  |   minProgrees ← 1
3  |   for i ← 1 to N do
4  |   |   if p_i ≤ minProgrees then
5  |   |   |   childStatus[i] ← children[i].Tick()
6  |   |   |   minProgrees ← min(minProgrees, p_i)
7  |   |   else
8  |   |   |   children[i].Pause()
9  |   if Σ_{i:childStatus[i]=Success} 1 ≥ M then
10 |   |   return Success
11 |   else if Σ_{i:childStatus[i]=Failure} 1 > N − M then
12 |   |   for i ← 1 to N do
13 |   |   |   children[i].Halt()
14 |   |   return Failure
15 |   return Running
16 Function Halt()
17 |   for i ← 1 to N do
18 |   |   children[i].Halt()
19 Function Pause()
20 |   for i ← 1 to N do
21 |   |   children[i].Pause()
```

---

for another one (*pause*) (e.g. a walking action moves the robot on a stable position with both feet on the ground when it has to stop, whereas it slows down speed of the motors when waiting for another action). Algorithm 8 shows the pseudocode of a Synchronized Parallel operator that explicitly defines these two routines. The Mutually Exclusive Parallel operator can be extended similarly.

B. Resolution of Progress Function

It sometimes impossible to define the progress value of a BT Action with a closed-form expression. This could result in the manual definition of the progress with a coarse resolution, with fixed value on a large portion of the state space. This can jeopardize the synchronized execution of two CBTs (e.g. the extreme case is where an action sets its progress to 1 only when it is completed and 0 otherwise).

C. Resource Assignment Policy

In Definition 14 and Algorithm 7, the user has to define a policy to access a critical resource. The policy has to implement aging to avoid starvation while it must give the resources to the current running action for enough time (e.g. a speech action must have a high chance to use the speaker while the robot is speaking). In a modular framework as BTs the choice of the priority value can require effort, since in some cases we must ensure that the priority of a node is the highest one. However, task scheduling algorithms provide effective solutions.

VIII. Use Cases Examples

In this section, we show two use cases in realistic scenarios where the problems formulated apply. One case involves the Syncronized Parallel operator and the other involves the Mutually Exclusive Parallel operator.

In both cases we use the BT in Figure 2. The action *Navigate* makes the robot follow a collision-free path, If such path exists. The action *Ask People to Move* uses the speaker to ask to make way. The action *Seek* makes the robot move its head looking for objects. Whenever an object is found, the robot stops moving the head until the object is correctly recognized.


Fig. 2. BT of Cases 1 and 2.

Case 1 (Action Synchronization): This case is an implementation of Example 1. When the BT starts its execution, both actions are executed and their progress advances as expected. When the robot finds an object on the floor, it takes time to recognized the object. This blocks the progress of the action *Sweep*. As a consequence, the robot stops navigating because the action *Navigate* no longer receive Ticks. When the object is recognized, the progress of Seek restarts, making the robot resume its navigation. Similarly, when the robot has to avoid obstacles, the Navigate action progresses slowly, as a consequence, the progress of the action *Seek* slows down.

Case 2 (Exclusive Use of Resources): This case is an implementation of Example 2. To reach an employee's desk, the robot enters a dark office. In this situation the condition *Light Level Good* returns false, making the ticks reach the action *Ask For Better Light*, which asks to turn on lights four times. An employee hears the robot's request and moves to the switch to turn on the lights. While doing so the employee moves in front of the robot obstructing its path. The robot is now unable to find another path and the condition *No People in Front* returns Failure. In this case, the action *Ask People to Move* cannot be executed as it requires the speaker to be available but it is used by the action *Ask For Better Light*. After a while the situation remains unchanged, the room is still dark and the employee is still in front of the robot. The priority of the action *Ask People to Move* becomes the highest one. The BT stops sending ticks to *Ask For Better Light* and starts sending ticks to *Ask People to Move*.

In both use cases the environment can be dynamic and unpredictable, a case where BTs shows advantages over other classical control architecture [14].

IX. Conclusions

In this paper, we proposed CBTs a generalization of BTs that allows the definition of two new control flow nodes for data and progress synchronization. We defined CBTs and the two new operators according to the state space formulation of BTs, in line with [14], analyzed the approach from a theoretical standpoint and provided use cases. This is the first step towards the synchronization of concurrent BTs.

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] I. Millington and J. Funge, *Artificial intelligence for games*. CRC Press, 2009.

[5] S. Rabin, *Game AI Pro*. CRC Press, 2014, ch. 6. The Behavior Tree Starter Kit.

[6] 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.

[7] K. R. Guerin, C. Lea, C. Paxton, and G. D. Hager, "A framework for end-user instruction of a robot assistant for manufacturing," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.

[8] D. Hu, Y. Gong, B. Hannaford, and E. J. Seibel, "Semi-autonomous simulated brain tumor ablation with raven ii surgical robot using behavior tree," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.

[9] A. Klökner, "Interfacing Behavior Trees with the World Using Description Logic," in *AIAA conference on Guidance, Navigation and Control, Boston*, 2013.

[10] P. Ögren, "Increasing Modularity of UAV Control Systems using Computer Game Behavior Trees," in *AIAA Guidance, Navigation and Control Conference, Minneapolis, MN*, 2012.

[11] 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.

[12] R. Brooks, "A robust layered control system for a mobile robot," *IEEE journal on robotics and automation*, vol. 2, no. 1, pp. 14–23, 1986.

[13] N. Nilsson, "Teleo-reactive programs for agent control," *Journal of artificial intelligence research*, 1994.

[14] M. Colledanchise and P. Ögren, *Behavior Trees in Robotics and AI: An Introduction*, ser. Chapman and Hall/CRC Artificial Intelligence and Robotics Series. Taylor & Francis Group, 2018.

[15] M. Colledanchise, A. Marzinotto, D. V. Dimarogonas, and P. Oegren, "The advantages of using behavior trees in multi-robot systems," in *ISR 2016: 47st International Symposium on Robotics; Proceedings of*. VDE, 2016, pp. 1–8.

[16] A. Champandard, "Enabling concurrency in your behavior hierarchy," *AIGameDev. com*, 2007.

[17] B. G. Weber, P. Mawhorter, M. Mateas, and A. Jhala, "Reactive planning idioms for multi-scale game ai," in *Computational Intelligence and Games (CIG), 2010 IEEE Symposium on*. IEEE, 2010, pp. 115–122.

[18] 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, April 2017.

[19] B. Hannaford, "Behavior Trees as a Representation for Medical Procedures," *ArXiv e-prints*, Jan. 2018.

[20] M. Mateas and A. Stern, "A behavior language for story-based believable agents," *IEEE Intelligent Systems*, vol. 17, no. 4, pp. 39–47, 2002.

[21] G. Taubenfeld, *Synchronization algorithms and concurrent programming*. Pearson Education, 2006.