

Verifying Safety of Behaviour Trees in Event-B

Matteo Tadiello

Elena Troubitsyna

KTH Royal Institute of Technology
Stockholm, Sweden

tadiello@kth.se

elenatro@kth.se

Behavior Trees (BT) are becoming increasingly popular in the robotics community. The BT tool is well suited for decision-making applications allowing a robot to perform complex behavior while being explainable to humans as well. Verifying that BTs used are well constructed with respect to safety and reliability requirements is essential, especially for robots operating in critical environments. In this work, we propose a formal specification of Behavior Trees and a methodology to prove invariants of already used trees, while keeping the complexity of the formalization of the tree simple for the final user. Allowing the possibility to test the particular instance of the behavior tree without the necessity to know the more abstract levels of the formalization.

1 Introduction

Autonomous Systems (AS) like Humanoid Robots, Autonomous Vehicles, or Unmanned Aerial Vehicles are becoming increasingly complex and need to interact with dynamic environments and with each other. For this reason, robots require tools to enable advanced perception and understanding of the environment, or capabilities to operate in complex situations. Artificial Intelligence is extending the capability of perception and action of the agents and allows robots to operate in environments not suitable for robots just a few years ago.

In most common scenarios the complexity of the environment requires to the robot to have different skills, the capability of different actions, and hence also a certain degree of reasoning and understanding of which action to take and when. A relevant example could be an urban road, with car, pedestrian, and signals. That's why planning becomes more important with the increasing complexity of the Operational Design Domain (ODD). Various techniques have been developed to solve this challenge. They include a variety of machine learning approaches as well as Behaviour Trees. Behaviour Trees started as a tool for video-game development of Non-Playable Character [19], and are now being applied also to robotics due to their flexibility and high usability.

Usually, robotics systems should be dependable, i.e., being able to operate without putting at risk others, or themselves, while continuing to provide their service in a reliable way. Formal approaches can be used to provide these characteristics and ensure the safety and reliability of the systems. Applying formal methods to AI techniques, however, is still challenging given the enormous search space and states possibility, and became particularly relevant when dealing with Neural Networks with billions of parameters. The importance of safety in robots that need to operate in a day-to-day basis in a critical environment is extremely important, as well as the capability of the system to explain why certain decisions have been made. For this reason Behaviour Trees (BT) have become particularly interesting in the robotics industry. This mathematical model follows simple principles and allow readability, modularity and re-usability. And for this reason, it has gain popularity especially in decision making tasks.

However, even with these properties BT can become quite complex and easy to contain design mistakes that could mine the safety and reliability of the system. For these reasons, it is important to ensure

M. Luckcuck and M. Farrell; M. Autili, L. Berardinelli,
A. Bucaioni and C. Pompilio (Eds.):

Formal Methods for Autonomous Systems and
Automated and verifiable Software sYstem DEvelopment
EPTCS 371, 2022, pp. 139–155, doi:10.4204/EPTCS.371.10

This work is dedicated to the public domain.

that certain properties are kept during the execution of the BT that ensure the dependability of the system built. The usage of formal methods to ensure the correctness of BT is still scarce, and difficult to build.

In this work, we try to advance in this direction by providing a formal specification of the Behaviour Tree framework that will allow the verification of correctness of specific BT instances and holding of interesting invariant properties that the system should maintain. In particular, we are able to define safety invariants and prove that these hold for specific BTs, while maintaining the re-usability and modularity of the tree structure. The system modeling has been built using the Event-B modeling method [2] and the Rodin [1] platform as a support for the refinements and mathematical proofs.

The rest of this paper is structured in the following way: we first provide an introductory explanation of Behaviour Trees and Event-B in the Background section, we continue with the Method section presenting the BT formal specification, followed by a case study section where we show a simple BT instance and the method to convert a BT to a formal specification that can be checked. We will then move to the related work and end with the conclusion section.

2 Background

2.1 Behaviour Trees

At a high level, a Behaviour Tree is a mathematical model that structures the switching of different tasks, which a virtual agent or a robot can perform [9]. This model uses a rooted tree where every node has a single parent and no loops. The nodes in a BT can be subdivided into three categories:

- **root:** the root node. It does not have parents, and it normally has just one child.
- **internal node:** also known as control flow nodes, those are nodes that are at the intermediary level of the tree. They all have a parent and at least a child. They represent how to navigate the tree.
- **leaf node:** also known as execution nodes, these nodes have a parent but do not have a child. They represent the tasks that the agent needs to perform (like "go to point A") or conditions (like "battery > than 30%") that needs to be true before executing an action.

In practice, the BT is navigated from the root to the leaves to decide which task needs to be executed at every time. To navigate the tree a process called *ticking* is used. An example of a BT and its ticking process can be seen in Fig 1. The navigation starts by ticking the root and then recursively ticking the left-most child of the node. When a leaf node is ticked it immediately replies with the status and sends the signal back to its parents, then the continuation of the navigation is pursued by following the rules of the internal nodes. A node is executed only if it has been ticked. At each time step, the tree is explored and returns the execution status of the agent. The output value can be *SUCCESS*, *RUNNING*, or *FAILURE* which, respectively, represent if the behavior of the agent, is in a success state, in a failure, or if it is still running. We call this the status of the tree.

The status of the tree is due to the composition of the statuses of the nodes in each subtree analyzed, and each node follows different rules to provide its status.

More specifically, for the leaf nodes we can distinguish two kinds:

- **Conditions:** these nodes represent guards check, they are normally placed before tasks nodes, they can provide just *SUCCESS* or *FAILURE*, but not *RUNNING*)
- **Tasks or Actions:** these nodes represent the actual actions executed by the agent. Once ticked, they immediately return *SUCCESS*, *RUNNING*, or *FAILURE*.

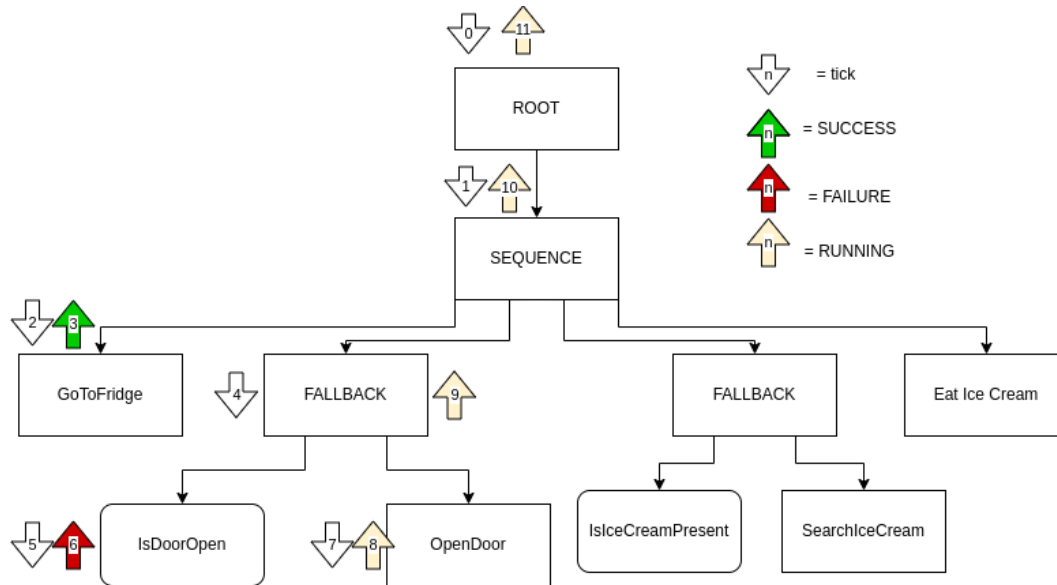


Figure 1: A sample of a BT and its ticking process. The execution of the tree returns RUNNING while executing the action of opening the door.

The leaf nodes are the ones that check the status of the agent and execute actions. The internal nodes instead define how leaf nodes are, or are not, executed.

The internal nodes, analyzed, are divided into two categories:

- **Sequence:** these nodes execute a series of children in sequence. The next child is executed only if the previous child has returned SUCCESS. The sequence node return to their parent SUCCESS if and only if all the children have returned SUCCESS. It returns RUNNING if one of the children is RUNNING, and FAILURE otherwise (hence one of the children has returned FAILURE). The sequence symbol is a box with the "→" label.
- **Fallback:** these nodes tick a series of children following a left-to-right order, starting from the leftmost child. The next child is executed only if the previous child has returned FAILURE. The fallback node return to their parent FAILURE if and only if all the children have returned FAILURE. It returns RUNNING if one of the children is RUNNING, and SUCCESS otherwise (hence one of the children has returned SUCCESS). The fallback symbol is a box with the "?" label.

BT can be extended also with other two internal nodes categories called PARALLEL and DECORATORS. Those nodes, less frequent, are not treated in this work.

The main strengths of Behaviour Tree are that they are simple and readable, and each subtree is reusable given the modularity property of this structure. But they can represent also complex behaviors for robots operating in challenging environments. Moreover, BT can also be extended with different properties like Stochastic BTs [5] or automatically generated as in [7].

2.2 Event-B and Rodin

The formal development framework used in this work is the *Event-B* formalism. In Event-B a system specification is defined using the concept of *abstract-state machine*. This model encapsulates a collection




Figure 2: Event-B machine and context

of variables and contains the set of operations to interact with the states of the machine called *Events*. We can thus define the behavior of the machine by describing the dynamic part of the state machine. Another component of the Event-B method is the so-called *context*, these components contain the static part of the system. For instance, in the context user-defined types are defined or constant components of the system. The context can be seen by the abstract machines as shown in 2.

The dynamic part of the system is represented by the variables of the abstract-state machines and the Cartesian product of their range represents the state space of the machine. The variables are strongly typed by invariant statements and are initialized using a special initialization event. The invariant statements, not only define the type of the variables but most importantly, are used to define properties that the system should hold for the entire time of execution. For instance, a possible invariant could be a safety property like: "the distance to the obstacles needs to be at least 5 meters from the robot".

The changing of the variables, and hence the state of the machine, is represented through a set of Event clauses defined in the machine. An event is normally defined as follows:

$$\mathbf{evt} \hat{=} \mathbf{any\ variables\ where\ guards\ then\ S\ end}$$

where *variables* is a new list of local variables, *guards* are predicates defining when the event can be triggered, and *S* is a set of statements or assignment that changes the state of the machine. However, both the variables and the guards are optional in the definition of an event. Multiple events can be enabled at the same time, in this case, any of them can be chosen to be executed in a non-deterministic way.

One of the strong points of Event-B is that it allows the possibility of machine refinement. This allows us to start the definition from a very high and abstract level, and include more details, variables, and events at every refinement. We can then introduce implementation details in steps while preserving functional correctness.

The semantics of an Event-B model is formulated as a collection of proof obligations [2], and the process of verification and generation and proving of proof obligations, has been facilitated by the Rodin platform [1].

3 Specification of Behaviour Trees in Event-B

The formal specification of a generic Behaviour Tree has been implemented and proven using Event-B and Rodin¹. To do that, we have used 4 different refinements. We built initially the abstract figure of the nodes, till arriving at the instances of the BT, which can be found as a refinement of the 4th level. Each abstract machine is coupled with a corresponding context. The four abstraction levels are described in detail in the following subsections, and are:

- Nodes
- Tree
- Behavior Tree
- Specific Instance of a BT.

3.1 Specifying a node

We start defining the nodes and certain properties that the nodes composing the Tree need to have. At this level, we are not interested in the connection and relationship between the nodes but just in their properties.

First of all, we define two new types to define the kind of node and the result that each node can return.

The *TYPE* type represents the category that each node can be and it is defined as follow:

$$TYPE = \{ROOT, SEQUENCE, FALLBACK, CONDITION, ACTION\}$$

It contains the types for the root, control flow nodes, and execution nodes, this will be used in a future refinement to define the different navigation behavior of the nodes.

The *RESULT* type, instead, represents the value of return that each node can contain. We define it as follows:

$$RESULT = \{SUCCESS, RUNNING, FAILURE, UNKNOWN\}$$

The first three elements are the normal return value of a node in a BT, the *UNKNOWN* element instead is used as the default value when a node has not been ticked yet, and hence the return value is unknown.

We wanted to define the single nodes as a data structure containing the following members:

```
node {
  int n_id; //constant value in context
  TYPE n_type; //constant value in context
  bool n_tick; //dynamic value in machine
  RESULT n_result; //dynamic value in machine
};
```

Where *n_id* represents a unique identifier of the node that starts from 0 with the root. This value is used to define an order on the nodes used in a future refinement to identify the order of the children, used to navigate the tree from left to right. And *n_tick* is a boolean variable to see if the node has been *ticked* or not.

Since in Event-B there is not a standard definition of a structure, the members of the nodes have been defined as relationships between the nodes and the corresponding type as follows:

¹The source code can be found here: https://github.com/tadteo/BT_eventb

$$\begin{array}{l}
n_type \in NODES \rightarrow TYPE \\
n_id \in NODES \rightarrow \mathbb{N} \\
n_tick \in NODES \rightarrow BOOL \\
n_result \in NODES \rightarrow RESULT
\end{array}$$

and initialized as:

$$\begin{array}{l}
n_tick \in NODES \times \{FALSE\} \\
n_result \in NODES \times \{UNKNOWN\}
\end{array}$$

where $NODES$ is the set of all possible nodes that can be used to build the tree, and that has been defined in the node context together with the types, and the constant values. n_id and n_type does not require initialization since they are defined in the context and remain constant.

3.2 Formal model of tree

In the first refinement, we give the nodes a structure and build the rules to construct a tree. We used the following four requirements to define the tree structure:

- **Req1:** There is a single root object
- **Req2:** Each object other than the root has a parent
- **Req3:** There are no loops in the parent structure
- **Req4:** Each object is reachable from the root

First of all, we define the *nodes* constant which represents the set of nodes (a subset of $NODES$) in the tree, with *root* as the $ROOT$ node.

For **Req1** we define an axiom that states that for every pair of nodes, if one is a $ROOT$ then the other is not $ROOT$:

$$\forall n1, n2. (n1 \in nodes \wedge n2 \in nodes \wedge n1 \neq n2 \wedge n_type(n1) = ROOT \implies n_type(n2) \neq ROOT) \quad (1)$$

For **Req2** we define the relation *parent* as following:

$$parent \in nodes \setminus \{root\} \rightarrow nodes \quad (2)$$

For the no loop property we use the inverse of the parent function (i.e. the children)² in the following way:

$$\forall n. (n \subseteq parent^{-1}[n] \implies n = \emptyset) \quad (3)$$

For **Req4** we use a transitive closure on a node relation, that in a graph represents a relation to any other node in the graph [14].

To achieve that, we first define a generic relation between two nodes, and then we define the transitive closure of it.

$$\begin{array}{l}
node_rel = NODES \leftrightarrow NODES \\
tcl \in node_rel \rightarrow node_rel \\
\forall r. (r \in node_rel \implies r \subseteq tcl(r)) \\
\forall r. (r \in node_rel \implies r; tcl(r) \subseteq tcl(r) // \text{The unfolding of tcl is also part of tcl}) \\
\forall r, t. (r \in node_rel \wedge r \subseteq t \wedge r; t \subseteq t \implies tcl(r) \subseteq t // \text{tcl(r) is least}) \\
\forall r. (r \in node_rel \implies r \cup (r; tcl(r)))
\end{array}$$

²In Event-B in Rodin, the inverse is written using the \sim symbol. In this case, the parent function would be written as: $parent \sim [\{node\}]$

3.3 Behavior Tree level

The Behavior Tree level starts from the definition of the tree present in the level above and adds all the rules to allow the correct navigation and return of the result.

In this level, we simulate the ticking process, which starts from a completely unvisited and unticked tree. As the first step, the root needs to be ticked. After this, a chain reaction of events brings to the exploration of the tree which eventually returns a result. The exploration ends whenever the result arrives at the root node. At this point, the tree got reinitialized and the root get ticked again, restarting the cycle.

To achieve correct navigation of the tree, each type of node needs to be treated separately with specific events.

3.3.1 Root related events

To deal with the root we have 3 distinct events:

The first **tick_root** is used to tick the root node. It is used whenever the tick status of the root is **FALSE** and the result is **UNKNOWN**:

```

tick_root ≐
  any node
  where
    n_tick(node) = FALSE
    n_type(node) = ROOT
    n_result(node) = UNKNOWN
    node ∈ NODES
    node ∈ nodes
  then
    n_tick(node) := TRUE

```

We then have an event that is used to tick the first unticked child of the root as soon as the root is ticked. In our specification, the root has just one child and is treated as a special node. This allows a simpler way to manage the start and end of an exploration of the tree.

```

root_ticked ≐
  any
    node
    child
  where
    n_tick(node) = TRUE
    n_type(node) = ROOT
    n_tick(child) = FALSE
    node ∈ nodes
    child ∈ {x | x ∈ parent-1{node} ∧ n_tick(x) = FALSE ∧
    ∀y · (y ∈ parent-1{node} ∧ n_tick(y) = FALSE ⇒ n_id(y) ≥ n_id(x)}
  then
    n_tick(child) := TRUE
    analyzing_subtree(child) := TRUE

```

Figure 3: The event root_ticked

```

root_reinitialize  $\hat{=}$ 
any
  node
where
   $n\_tick(node) = TRUE$ 
   $n\_type(node) = ROOT$ 
   $n\_result(node) \neq UNKNOWN$ 
   $node \in NODES$ 
   $node \in nodes$ 
then
   $n\_result := NODES \times \{UNKNOWN\}$ 
   $n\_tick := NODES \times \{FALSE\}$ 

```

Figure 4: The event root_reinitialize

To identify the first unticked child we make use of the id of the node. The id of the nodes has been given in a way that respects the order of visiting of the tree, following a left-to-right order, with siblings having the leftmost with a smaller id and the rightmost with the highest id. When choosing which node to visit, we select as a child the node that has the n_id smaller than all the other nodes, and that is not yet ticked. Finally, when the ticking process return to the root an event called **result_arrived** is used to copy the result of the child of the root to the root.

```

result_arrived  $\hat{=}$ 
any
  node
  child
where
   $node \in NODES$ 
   $n\_tick(node) = TRUE$ 
   $n\_type(node) = ROOT$ 
   $n\_result(node) = UNKNOWN$ 
   $child \in parent^{-1}[\{node\}]$ 
   $child \in \{x | x \in parent^{-1}[\{node\}] \wedge n\_result(x) \neq UNKNOWN\}$ 
then
   $n\_result(node) := n\_result(child)$ 

```

Before restarting the ticking process an event called **root_reinitialize** is used to reinitialize the tree nodes as unexplored and with unknown results.

3.3.2 Leaf nodes related events

The leaf nodes can be of two types: *CONDITION* or *ACTION*. The leaf nodes are the only ones that directly provide a result with the ticking of the nodes. However, these events need to be refined at the next level for each case of result (*SUCCESS*, *FAILURE* for the *CONDITION* nodes and *SUCCESS*, *FAILURE* and *RUNNING* for the *ACTION* nodes). And providing the correct logic to return the result wanted, since for each BT the logic to return the result can be different.

At this level of abstraction, the only action executed by these nodes is setting the *analyzing_subtree* variable of the parent node to FALSE. This is done as follows:

$$analyzing_subtree(parent(node)) := FALSE \quad (4)$$

3.3.3 Internal nodes related events

The internal nodes are the ones in charge of controlling the flow of execution and analysis of the tree. In this work, we have analyzed the **sequence** node and the **fallback** node. To formalize the behavior of this kind of node, we have used 5 different events for each kind covering all the possible ways in which the node can be found while analyzing the tree. At each step of the tree exploration, the control flow nodes either have enough information to return a value to the parents or need to continue the analysis of the children and tick the next child on the list. The five events have the following form: 1. **_ticked_initial*, 2. **_ticked_success*, 3. **_ticked_failure*, 4. **_ticked_running*, 5. **_ticked_continue*, with *** used as a wildcard that can be substituted with the name of the node (like *fallback_ticked_continue*).

The **_ticked_initial* event is used to start the ticking process of the children and it's only purpose is to tick the first children of these nodes.

Let's analyze now the events of the fallback node.

The definition is "*The Fallback node [...] corresponds to routing the ticks to its children from the left until it finds a child that returns either SUCCESS or RUNNING, then it returns SUCCESS or RUNNING accordingly to its own parent. It returns FAILURE if and only if all its children return FAILURE. Note that when a child returns RUNNING or SUCCESS, the Fallback node does not route the ticks to the next child (if any)*" [9].

To be able to route the ticks to the correct child (from left to right), we first need to correctly identify the first unticked child and last analyzed child. To do that we use a similar approach to the one used in the *root_ticked* event. We use the inverse of *parent* to identify the children of a node and then we use the *n_id* of the children to analyze the order, finally, we use the *n_result* and *n_tick* variables to check the correct value we are looking for.

For instance, to identify if the last child analyzed has returned the *RUNNING* result (used in *fallback_ticked_running*, Figure 5, we use the following guard:

$$\begin{aligned} \exists x \cdot (x \in parent^{-1}[\{node\}] \wedge n_tick(x) = TRUE \wedge \\ \forall y \cdot (y \in parent^{-1}[\{node\}] \wedge n_tick(y) = TRUE \implies n_id(y) \leq n_id(x)) \\ \wedge n_result(x) = RUNNING) \quad (5) \end{aligned}$$

In this code, we check if exists a node of the children set, which has been ticked, that has the greatest id of all the other ticked children, and that has the result set to *RUNNING*. To find the first unticked children instead we use the same approach but instead check that the *n_tick* variable is set to false and that the *n_id* value is the minimum of the unticked children. This can be seen in the events *fallback_ticked_continue* Fig. 5, and *sequence_ticked_continue* Fig. 6,

In some events, we want also to be able to identify whenever all the children have been analyzed (ticked). In these cases, we just check that the set of children of a node that has *n_tick* set to FALSE is an empty set. The guard used to check this is:

$$\{x \mid x \in parent^{-1}[\{node\}] \wedge n_tick(x) = FALSE\} = \emptyset \quad (6)$$

With these guards, we are now able to define the 4 events that allow us to correctly manage the Fallback node (Figure 5).

Events	
<pre> fallback_ticked_running ≐ any node where n_type(node) = FALLBACK n_tick(node) = TRUE n_result(node) = UNKNOWN ∃x · (x ∈ parent⁻¹[{node}] ∧ n_tick(x) = TRUE ∧ ∀y · (y ∈ parent⁻¹[{node}] ∧ n_tick(y) = TRUE ⇒ n_id(y) ≤ n_id(x)) ∧ n_result(x) = RUNNING) then n_result(node) := RUNNING analyzing_subtree(parent(node)) := FALSE end fallback_ticked_failure ≐ any node where n_type(node) = FALLBACK n_tick(node) = TRUE n_result(node) = UNKNOWN ∃x · (x ∈ parent⁻¹[{node}] ∧ n_tick(x) = TRUE ∧ ∀y · (y ∈ parent⁻¹[{node}] ∧ n_tick(y) = TRUE ⇒ n_id(y) ≤ n_id(x)) ∧ n_result(x) = FAILURE) {x x ∈ parent⁻¹[{node}] ∧ n_tick(x) = FALSE} = ∅ then n_result(node) := FAILURE analyzing_subtree(parent(node)) := FALSE end </pre>	<pre> fallback_ticked_success ≐ any node where n_type(node) = FALLBACK n_tick(node) = TRUE n_result(node) = UNKNOWN ∃x · (x ∈ parent⁻¹[{node}] ∧ n_tick(x) = TRUE ∧ ∀y · (y ∈ parent⁻¹[{node}] ∧ n_tick(y) = TRUE ⇒ n_id(y) ≤ n_id(x)) ∧ n_result(x) = SUCCESS) then n_result(node) := SUCCESS analyzing_subtree(parent(node)) := FALSE end fallback_ticked_continue ≐ any node, child where n_type(node) = FALLBACK n_tick(node) = TRUE n_result(node) = UNKNOWN ∃x · (x ∈ parent⁻¹[{node}] ∧ n_tick(x) = TRUE ∧ ∀y · (y ∈ parent⁻¹[{node}] ∧ n_tick(y) = TRUE ⇒ n_id(y) ≤ n_id(x)) ∧ n_result(x) = FAILURE) {x x ∈ parent⁻¹[{node}] ∧ n_tick(x) = FALSE} = ∅ node ∈ nodes child ∈ {x ∈ parent⁻¹[{node}] ∧ n_tick(x) = FALSE} ∀y · (y ∈ parent⁻¹[{node}] ∧ n_tick(y) = FALSE ⇒ n_id(y) ≥ n_id(x)) analyzing_subtree(node) := FALSE then n_tick(child) := TRUE analyzing_subtree(node) := TRUE end end </pre>

Figure 5: The events to manage the Fallback node

At every event either we return a result or we tick a child node. And we set the *analyzing_subtree* label true or false accordingly.

For the **Sequence** node, we follow a similar approach we the modified event to control it according to its definition.

The definition is "The Sequence node [...], corresponds to routing the ticks to its children from the left until it finds a child that returns either Failure or Running, then it returns Failure or Running accordingly to its own parent. It returns success if and only if all its children return Success. Note that when a child returns Running or Failure, the Sequence node does not route the ticks to the next child (if any)." [9].

Events	
<pre> sequence_ticked_running $\hat{=}$ any node where n_type(node) = SEQUENCE n_tick(node) = TRUE n_result(node) = UNKNOWN $\exists x \cdot (x \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(x) = \text{TRUE} \wedge$ $\forall y \cdot (y \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(y) = \text{TRUE} \implies$ $n_id(y) \leq n_id(x) \wedge n_result(x) = \text{RUNNING})$ then n_result(node) := RUNNING analyzing_subtree(parent(node)) := FALSE end sequence_ticked_failure $\hat{=}$ any node where n_type(node) = SEQUENCE n_tick(node) = TRUE n_result(node) = UNKNOWN $\exists x \cdot (x \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(x) = \text{TRUE} \wedge$ $\forall y \cdot (y \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(y) = \text{TRUE} \implies$ $n_id(y) \leq n_id(x) \wedge n_result(x) = \text{FAILURE})$ then n_result(node) := FAILURE analyzing_subtree(parent(node)) := FALSE end </pre>	<pre> sequence_ticked_success $\hat{=}$ any node where n_type(node) = SEQUENCE n_tick(node) = TRUE n_result(node) = UNKNOWN $\exists x \cdot (x \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(x) = \text{TRUE} \wedge$ $\forall y \cdot (y \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(y) = \text{TRUE} \implies$ $n_id(y) \leq n_id(x) \wedge n_result(x) = \text{SUCCESS})$ $\{x \mid x \in \text{parent}^{-1}[\{\text{node}\}]$ $\wedge n_tick(x) = \text{FALSE}\} = \emptyset$ node \in nodes then n_result(node) := SUCCESS analyzing_subtree(parent(node)) := FALSE end sequence_ticked_continue $\hat{=}$ any node, child where n_type(node) = SEQUENCE n_tick(node) = TRUE n_result(node) = UNKNOWN $\exists x \cdot (x \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(x) = \text{TRUE} \wedge$ $\forall y \cdot (y \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(y) = \text{TRUE} \implies$ $n_id(y) \leq n_id(x) \wedge n_result(x) = \text{SUCCESS})$ $\{x \mid x \in \text{parent}^{-1}[\{\text{node}\}]$ $\wedge n_tick(x) = \text{FALSE}\} = \emptyset$ node \in nodes child \in $\{x \in \text{parent}^{-1}[\{\text{node}\}]\}$ $\wedge n_tick(x) = \text{FALSE} \wedge$ $\forall y \cdot (y \in \text{parent}^{-1}[\{\text{node}\}] \wedge n_tick(y) = \text{FALSE} \implies$ $n_id(y) \geq n_id(x))$ analyzing_subtree(node) := FALSE then n_tick(child) := TRUE analyzing_subtree(node) := TRUE end end </pre>

Figure 6: The events to manage the Sequence node

The guards and events are similar to the Fallback node but with different logic. The events of the Sequence node can be seen in Figure 6.

3.4 Instantiating generic specification

The fourth refinement finally allows us to deal with real instances of BTs. At this level, actual Behaviour Trees can be verified and checked for invariants of interest for the agent. We now explain how the context and the machine need to be built to properly check an existing tree.

The context contains the static part of the machine, the BT in itself is a static part of the agent, and for this reason, the definition of the topology and of the nodes type of the BT needs to be defined here.

To define a BT instance we start by extending the *BT_level* and listing all the nodes contained in the tree (excluding the root, already defined in the previous refinement), as constants.

We then need to define the axioms regarding the type: that every node defined in the constant part is of type *nodes*. After this we need to initialize all the nodes giving to them a *n_type*, and a *n_id*. Finally, we construct the parent function, defining for each node which node is its parent, defining in this way the topology of the tree.

Numbering the nodes correctly is essential to correctly navigate a tree. To correctly give a *n_id* to the nodes, is required to follow the same principle as a breadth-first search: firstly considering the depth, and then starting counting the children from left to right. Therefore, The root, receive the number 0. We will then give to all its children the next numbers starting by counting from the left-most node (i.e. the first to be ticked). When all the children have been assigned a number we start to analyze the children of the first child, numbering them left-to-right in the same way as before. And so on, until all the nodes have been numbered.

Moving on the machine of the BT instance, we start as a refinement of the *BT_machine* of the previous level. The only effort required here is the extension and refinement of the events of the leaf nodes. Since the leaf nodes (conditions and actions) are very specific to the application, we need to extend the corresponding events for any action and condition we want to check. For instance, let's suppose that our agent has two condition nodes: *condition_1* and *condition_2*, the first condition checks if *the agent is more than 3 meters from an obstacle*, and the second that *the battery level is more than 30%*. To cover the two conditions we would need to extend the *condition_ticked* event of the BT level four times: 1. *condition_1_ticked_success*, 2. *condition_1_ticked_failure*, 3. *condition_2_ticked_success*, 4. *condition_2_ticked_failure*, covering all the possible results that the nodes can output. The modifications to be added would be minimal, in the *condition_1* case we would just need a guard to check the condition (something like $distance_to_obstacle \geq 3$) and an action that explain which value to return ($n_result(node) := SUCCESS$). For the failure case, we would add the condition for failure ($distance_to_obstacle < 3$) and the corresponding returning value ($n_result(node) := FAILURE$)

Moreover, if we want to check some safety invariant or other parameters, we would need to model the part of the environment we are interested to check and how the various actions would change the environment itself. For instance, if a robot is moving in an environment with some obstacles we could define two variables (*pos_x*, *pos_y* to define the position of the robot and another variable (*distance_to_obstacle*), to keep track of the distance to the closest obstacle. These would then be modified by actions. For instance, if we have an action "move one step up", in the event of the corresponding action we would also need to add something like: $pos_y := pos_y + 1$.

4 Case Study – A moving robot

As a case study, we analyzed a very simple case, considering a robot that is moving toward a wall. The goal is to avoid the robot from going too close to a certain distance to the obstacle Fig: 7.




Figure 7: Case study of a robot moving towards a wall, without going too close to it. The robot can move closer to the wall, but it is not allowed to enter the unsafe area.




Figure 8: Behaviour tree used to model the decision-making process of the agent.

To model this simple environment we have defined in the machine different variables: 1. *distance_to_object*, 2. *time*, 3. *prev_time*. We use the *distance_to_object* variable to identify the position of the robot and hence its distance to the wall, while a *time* variable (together with the previous time step) is used to control the real-time passed during the execution.

To define the behavior of the robot we have defined BT represented in Figure 8. As invariant to be checked we set:

$$distance_to_object \geq 3$$

In the initialization, we set the robot at a distance of 10 meters from the wall.

The first leaf node of the tree is a condition that the robot needs to be at least 5 meters from the wall. If this condition holds then the robot can perform the action of moving one step closer to the wall.

To model this environment, we followed the same approach described at the instance level above. We first defined the tree in the context described below. In the machine, we define the variables described above to define the environment and we refine the leaf node events.

For the condition, we create the two refinements of the event for the *SUCCESS* and *FAILURE* case as described in the methodology section. For the action, we simplify the scenario enforcing that, once selected, the action will take effect instantly and that can never fail. Allowing us to refine the action event just for the *SUCCESS* case.

```

context robot_wall_context
extends 03_BT_context

constants
  sequence_1
  condition_1
  action_1
axioms
  sequence_1 ∈ nodes
  condition_1 ∈ nodes
  action_1 ∈ nodes
  n_type = {root ↦ ROOT, sequence_1 ↦ SEQUENCE, condition_1 ↦ CONDITION, action_1 ↦ ACTION}
  n_id = {root ↦ 0, sequence_1 ↦ 1, condition_1 ↦ 2, action_1 ↦ 3}
  parent = {sequence_1 ↦ root, condition_1 ↦ sequence_1, action_1 ↦ sequence_1}

```

Figure 9: The context of the BT for the robot moving towards the wall

We can note that the action ticked event is very similar to the previous definition, with the only difference due to the addition of the last two propositions.

```

action_ticked_success ≐
REFINES action_ticked
any node
where
  n_type(node) = ACTION // not changed
  n_tick(node) = TRUE // not changed
  n_result(node) = UNKNOWN // not changed
  node ∈ nodes // not changed
then
  analyzing_subtree(parent(node)) := FALSE // not changed
  n_result(node) := SUCCESS // added in refinement
  distance_to_object := distance_to_object - 1 // added in refinement
end

```

Figure 10: The context of the BT for the robot moving towards the wall

We then added a statement to increase the timestep whenever the result reaches the root node, considering an entire exploration of the tree as a timestep in the real world since the changes in the position of the robot have been applied to the action event.

In general, just with these modifications at the machine from the BT level, we are able to ensure the safety properties of the system by proving the safety invariant that we have defined for this BT. This

ensures us that, after the initialization, the safety statement that we have described will hold for each possible state that the BT can reach. This can be expanded, with multiple safety statements and multiple requirements, allowing us to formally prove the safety of the system that will implement this specific BT.

5 Related Works

Although Behavior Trees have started to become widely used in various robotics applications, still little research has been done to increase the dependability of robots using them. The importance of providing safe and reliable agents becomes more and more critical, especially given the fact that these tools are starting to be used at a larger level, especially with the support of open source software like the Nav2 package for ROS [13]. Work in the safety assessment of Behavior Trees has been done, for instance, [12], but these studies normally refer to specific cases and become difficult to generalize for other BT. BT can also be safely generated as in [17, 6] and that BT is an optimal tool to build safe and reliable systems [8]. However, formal verification of safety properties is rarely done on new trees especially in manually constructed ones, risking in this way the overall dependability of the systems using them. Other works propose the formal verification of BT using Linear Temporal Logic [4], but they do not provide a tool to support the verification process.

Event-B was used for modelling robotic systems by Troubitsyna et al. [11, 18, 15]. These works adopted goal-oriented model as a basis of robot's decision making and demonstrated goal reachability despite robot's failures. However, these works rely on an assumption that the robot's operating environment remains static, i.e., once a certain goal is accomplished, its status remains unchanged throughout the entire system execution. In our work, reliance on BT allows us to adopt dynamic model of system environment, because at each tick the status of all tasks is getting checked, i.e., this work offers a more flexible basis for modeling robot's decision making.

6 Conclusion

In this work, we present a formal specification of Behaviour Trees in Event-B. We also provide a methodology to test new instances of BT in a fairly simple manner for newcomers hiding the definition of the BT at more abstract levels, allowing a user to focus just on test and checking the properties the specific BT needs to hold. Our approach aimed to facilitate the formal verification of Behaviour Trees without the need to build the formalization of the tree and the control flow node from scratch and simplifying also the modeling of the leaf nodes.

In future works, we propose to improve the actual specification by adding the support for parallel nodes and the most common decorator used in practice. Moreover, we would like to simplify further the possibility to check already used BT, creating a tool to automatically import trees built using already used libraries like `py_tree` [16] or `Behaviour_Tree.CPP` [3]. Of particular interest is also the analysis of already used BT in robotic frameworks like the ones used in the Navigation 2 [13] package in ROS.

References

- [1] *Event-B platform*,. Available at <http://www.event-b.org/>.
- [2] Jean-Raymond Abrial (2010): *Modeling in event-B: System and software engineering*. Cambridge University Press, doi:10.1017/CBO9781139195881.

- [3] BehaviorTree: *Behaviortree/BehaviorTree.CPP: Behavior trees library in C++. batteries included*. Available at <https://github.com/BehaviorTree/BehaviorTree.CPP>.
- [4] Oliver Biggar & Mohammad Zamani (2020): *A Framework for Formal Verification of Behavior Trees With Linear Temporal Logic*. *IEEE Robotics and Automation Letters* 5(2), pp. 2341–2348, doi:10.1109/LRA.2020.2970634. Conference Name: IEEE Robotics and Automation Letters.
- [5] Michele Colledanchise, Alejandro Marzinotto & Petter Ögren (2014): *Performance analysis of stochastic behavior trees*. pp. 3265–3272, doi:10.1109/ICRA.2014.6907328.
- [6] Michele Colledanchise, Richard M. Murray & Petter Ögren (2017): *Synthesis of Correct-by-Construction Behavior Trees*. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE Press, p. 6039–6046, doi:10.1109/IROS.2017.8206502.
- [7] Michele Colledanchise, Ramvijas Parasuraman & Petter Ögren (2015): *Learning of Behavior Trees for Autonomous Agents*. *CoRR* abs/1504.05811, doi:10.1109/TG.2018.2816806.
- [8] Michele Colledanchise & Petter Ögren (2014): *How Behavior Trees modularize robustness and safety in hybrid systems*. In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1482–1488, doi:10.1109/IROS.2014.6942752.
- [9] Michele Colledanchise & Petter Ögren (2018): *Behavior Trees in Robotics and AI: An Introduction*. CRC Press, doi:10.1201/9780429489105. Available at <http://arxiv.org/abs/1709.00084>.
- [10] Razan Ghzouli, Thorsten Berger, Einar Broch Johnsen, Swaib Dragule & Andrzej Wąsowski (2020): *Behavior trees in action: a study of robotics applications*. In: *Proceedings of the 13th ACM SIGPLAN International Conference on Software Language Engineering, SLE 2020*, Association for Computing Machinery, New York, NY, USA, pp. 196–209, doi:10.1145/3426425.3426942.
- [11] Linas Laibinis, Inna Pereverzeva & Elena Troubitsyna (2017): *Formal reasoning about resilient goal-oriented multi-agent systems*. *Science of Computer Programming* 148, pp. 66–87, doi:10.1016/j.scico.2017.05.008. Available at <https://www.sciencedirect.com/science/article/pii/S0167642317301181>. Special issue on Automated Verification of Critical Systems (AVoCS 2015).
- [12] Peter A. Lindsay, Kirsten Winter & Nisansala Yatapanage (2010): *Safety Assessment Using Behavior Trees and Model Checking*. In: *2010 8th IEEE International Conference on Software Engineering and Formal Methods*, IEEE, Pisa, Italy, pp. 181–190, doi:10.1109/SEFM.2010.23. Available at <http://ieeexplore.ieee.org/document/5637427/>.
- [13] Steve Macenski, Francisco Martín, Ruffin White & Jonatan Ginés Clavero (2020): *The Marathon 2: A Navigation System*. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2718–2725, doi:10.1109/IROS45743.2020.9341207. Available at <http://arxiv.org/abs/2003.00368>.
- [14] W.F. McColl & K. Noshita (1986): *On the number of edges in the transitive closure of a graph*. *Discrete Applied Mathematics* 15(1), pp. 67–73, doi:10.1016/0166-218X(86)90020-X. Available at <https://www.sciencedirect.com/science/article/pii/0166218X8690020X>.
- [15] Inna Pereverzeva, Elena Troubitsyna & Linas Laibinis (2012): *Formal Goal-Oriented Development of Resilient MAS in Event-B*. In Mats Brorsson & Luís Miguel Pinho, editors: *Reliable Software Technologies – Ada-Europe 2012*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 147–161. doi:10.1007/978-3-642-30598-6_11.
- [16] Splintered-Reality: *Splintered-reality/py_trees: Python implementation of behaviour trees*. Available at https://github.com/splintered-reality/py_trees.
- [17] Tadewos G. Tadewos, Laya Shamgah & Ali Karimoddini (2019): *Automatic Safe Behaviour Tree Synthesis for Autonomous Agents*. In: *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 2776–2781, doi:10.1109/CDC40024.2019.9030183.
- [18] Inna Vistbakka & Elena Troubitsyna (2021): *Modelling resilient collaborative multi-agent systems*. *Computing* 103, pp. 535–557, doi:10.1007/s00607-020-00861-2. Available at <https://link.springer.com/article/10.1007/s00607-020-00861-2>.

- [19] Michael Waltham & Deshen Moodley (2016): *An Analysis of Artificial Intelligence Techniques in Multiplayer Online Battle Arena Game Environments*. In: *Proceedings of the Annual Conference of the South African Institute of Computer Scientists and Information Technologists, SAICSIT '16*, Association for Computing Machinery, New York, NY, USA, doi:10.1145/2987491.2987513.