

Decentralised Evaluation of Temporal Patterns over Component-based Systems at Runtime^{*}

Olga Kouchnarenko^{1,2} and Jean-François Weber¹

¹ FEMTO-ST CNRS and University of Franche-Comté, Besançon, France

² Inria/Nancy-Grand Est, France

{okouchnarenko,jfweber}@femto-st.fr

Abstract. Self-adaptation allows systems to modify their structure and/or their behaviour depending on the environment and the system itself. Since reconfigurations must not happen at any but in suitable circumstances, guiding and controlling dynamic reconfigurations at runtime is an important issue. This paper contributes to two essential topics of the self-adaptation—a runtime temporal properties evaluation, and a decentralization of control loops. It extends the work on the adaptation of component-based systems at runtime via policies with temporal patterns by providing *a) a specific progressive semantics of temporal patterns and b) a decentralised method which is suitable to deal with temporal patterns of component-based systems at runtime.* The implementation with the *GROOVE* tool constitutes a practical contribution.

1 Introduction

Self-adaptation—the ability of systems to modify their structure and/or their behaviour in response to their interaction with the environment and the system itself, and their goals—is an important and active research field with applications in various domains [1]. Since dynamic reconfigurations that modify the architecture of component-based systems without incurring any system downtime must not happen at any but in suitable circumstances, adaptation policies are used to guide and control reconfigurations at runtime. For triggering adaptation policies and specifying behaviours of component-based systems, a linear temporal logic based on Dwyer’s work on patterns and scopes [2], called FTPL, has been used in [3]. In this adaptation context, choosing a suitable adaptation policy in a current component-based system configuration depends on a *runtime temporal patterns evaluation* which is one of the essential topics of the self-adaptation [1].

We consider *open* component-based systems interacting with their environment, therefore, their behaviour depends on both external and internal events. Since our component-based systems are modelled by infinite state transition systems, for our pattern-based verification to remain tractable, we consider a non-blocking environment with incomplete information about the component-based system that enables all the external events.

^{*} This work has been partially funded by the Labex ACTION, ANR-11-LABX-0001-01.

In this setting, providing values for temporal patterns is a difficult task. In [3], a centralised evaluation of temporal patterns at runtime has been proposed. In order to contribute to *decentralization of control loops*—another self-adaptation topic, this paper addresses the FTPL *decentralised* evaluation problem on a re-configuration path, and presents a method that is suitable to deal with temporal patterns of component-based systems. Indeed, as these patterns contain conjunctions or disjunctions of properties over components’ parameters and relations, the evaluation of temporal patterns in a decentralised manner makes sense, and the sooner the better.

Inspired by the LTL decentralised evaluation [4] for closed systems, this paper provides a progressive FTPL semantics allowing a decentralised evaluation of FTPL formulae over open component-based systems – the first contribution. The second contribution consists of an algorithm to answer the temporal pattern decentralised evaluation problem in \mathbb{B}_4 and of the correctness and uniqueness results saying that whenever an FTPL property is evaluated in the decentralised manner, it matches the FTPL evaluation using the basic semantics in [3]. The implementation with the *GROOVE* tool [5] constitutes a practical contribution.

Related work. When checking properties of open systems, the idea is to satisfy a property no matter how the environment behaves. For non-universal temporal logics, this problem, called *module-checking*, is in general much harder than model-checking of closed systems in finite as well as in infinite settings [6,7], and it becomes undecidable with imperfect information about the control states [8]. Fortunately, for universal temporal logics as LTL, the module checking problem with both complete or incomplete information remains decidable in finite setting [6]; in particular, it is PSPACE-complete for LTL.

As temporal properties often cannot be evaluated to true or false during the system execution, in addition to *true* and *false* values, *potential true* and *potential false* values are used whenever an observed behaviour has not yet led to an acceptance or a violation of the property under consideration, forming the \mathbb{B}_4 domain like in RV-LTL [9]. Like in [10], in our framework external events can be seen as invocations of methods performed by (external) sensors when a change is detected in their environment.

Let us remark that this work is motivated by applications in numerous frameworks that support the development of components together with their monitors/controllers, as, e.g., Fractal [11], CSP||B [12], FraSCAti [13], etc.

More generally, this paper aims to contribute to the development of new verification approaches for complex systems that integrate ideas of distributed algorithms [14].

Layout of the paper. After a short overview of a component-based model and of a linear temporal patterns logic in Section 2, Section 3 presents a specific progressive semantics of temporal patterns. Afterwards, Section 4 addresses the temporal pattern decentralised evaluation problem on a reconfiguration path by providing an algorithm for such an evaluation in \mathbb{B}_4 . Section 5 describes the implementation with the *GROOVE* tool and details an example of a location composite component. Finally, Section 6 presents our conclusion.

2 Background: Reconfiguration Model and Temporal Patterns

In this section we recall the component-based model and the temporal patterns introduced, respectively, in [15] and [16].

2.1 Reconfiguration Model

The reconfigurations we consider here make the component-based architecture evolve dynamically. They are combinations of basic reconfiguration operations such as instantiation/destruction of components; addition/removal of components; binding/unbinding of component interfaces; starting/stopping components; setting parameter values of components. In the remainder of the paper, we focus on reconfigurations that are combinations of basic operations.

In general, a system configuration is the specific definition of the elements that define or prescribe what a system is composed of. As in [15], we define below a configuration to be a set of architectural elements (components, required or provided interfaces, and parameters) together with relations (binding, delegation, state, etc.) to structure and to link them, as depicted in Fig. 1.

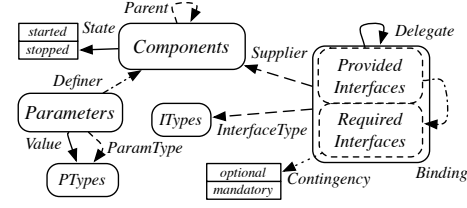


Fig. 1: Configurations = architectural elements and relations

Definition 1 (Configuration). A configuration c is a tuple $\langle Elem, Rel \rangle$ where

- $Elem = Components \uplus Interfaces \uplus Parameters \uplus Types$ is a set of architectural elements, such that
 - $Components$ is a non-empty set of the core entities, i.e components;
 - $Interfaces = ReqInterfaces \uplus ProvInterfaces$ is a finite set of the (required and provided) interfaces;
 - $Parameters$ is a finite set of component parameters;
 - $Types = ITypes \uplus PTypes$ is a finite set of the interface types and the parameter data types;
- $Rel = \left\{ \begin{array}{l} Container \uplus ContainerType \uplus Parent \\ \uplus Binding \uplus Delegate \uplus State \uplus Value \end{array} \right.$ is a set of architectural relations which link architectural elements, such that
 - $Container : Interfaces \uplus Parameters \rightarrow Components$ is a total function giving the component that supplies the considered interface or the component of a considered parameter;
 - $ContainerType : Interfaces \uplus Parameters \rightarrow Types$ is a total function that associates a type with each required/provided interface, or with a parameter;

- $Parent \subseteq Components \times Components$ is a relation linking a sub-component to the corresponding composite component³;
- $Binding : ProvInterfaces \rightarrow ReqInterfaces$ is a partial function that binds together a provided interface and a required one;
- $Delegate : Interfaces \rightarrow Interfaces$ is a partial function that expresses delegation links;
- $State : Components \rightarrow \{started, stopped\}$ is a total function giving the status of instantiated components;
- $Contingency : ReqInterfaces \rightarrow \{mandatory, optional\}$ is a total function to characterise the required interfaces;
- $Value : Parameters \rightarrow \bigcup_{ptype \in PType} ptype$ is a total function that gives the current value of each parameter.

Given a set of configurations $\mathcal{C} = \{c, c_1, c_2, \dots\}$, let CP be a set of configuration properties on the architectural elements and the relations between them specified using first-order logic formulae. The interpretation of functions, relations, and predicates is done according to basic definitions in [17] and in [15]. A configuration *interpretation* function $l : \mathcal{C} \rightarrow CP$ gives the largest conjunction of $cp \in CP$ evaluated to true on $c \in \mathcal{C}$ ⁴.

Among all the configuration properties, the architectural *consistency constraints* CC express requirements on component assembly common to all the component architectures. These constraints introduced in [18] allow the definition of *consistent configurations* regarding, in particular, the following rules:

- a component *supplies* one provided interface, at least;
- the composite components have no parameter;
- a sub-component must not include its own parent component;
- two bound interfaces must have the same interface type and their containers are sub-components of the same composite;
- when binding two interfaces, there is a need to ensure that they have not been involved in a delegation yet; similarly, when establishing a delegation link between two interfaces, the specifier must ensure that they have not yet been involved in a binding;
- a provided (resp. required) interface of a sub-component is delegated to at most one provided (resp. required) interface of its parent component; the interfaces involved in the delegation must have the same interface type;
- a component is *started* only if its mandatory required interfaces are bound or delegated.

Definition 2 (Consistent configuration). Let $c = \langle Elem, Rel \rangle$ be a configuration and CC the architectural consistency constraints. The configuration c is consistent, written *consistent*(c), if $l(c) \Rightarrow CC$.

Let \mathcal{R} be a finite set of reconfiguration operations, and *run* a generic running operation. The possible evolutions of the component architecture via the reconfiguration operations are defined as a transition system over $\mathcal{R}_{run} = \mathcal{R} \cup \{run\}$.

³ For any $(p, q) \in Parent$, we say that q has a sub-component p , i.e. p is a child of q .

⁴ By definition in [17], this conjunction is in CP .

Definition 3 (Reconfiguration model). *The operational semantics of component systems with reconfigurations is defined by the labelled transition system $S = \langle \mathcal{C}, \mathcal{C}^0, \mathcal{R}_{run}, \rightarrow, l \rangle$ where $\mathcal{C} = \{c, c_1, c_2, \dots\}$ is a set of consistent configurations, $\mathcal{C}^0 \subseteq \mathcal{C}$ is a set of initial configurations, \mathcal{R} is a finite set of reconfigurations, $\rightarrow \subseteq \mathcal{C} \times \mathcal{R}_{run} \times \mathcal{C}$ is the reconfiguration relation.*

Let us write $c \xrightarrow{ope} c'$ when c' is reached from c by $ope \in \mathcal{R}_{run}$. An evolution path σ (or a path for short) in S is a (possibly infinite) sequence of configurations c_0, c_1, c_2, \dots such that $\forall i \geq 0. (\exists ope_i \in \mathcal{R}_{run}. (c_i \xrightarrow{ope_i} c_{i+1}))$. Let $\sigma(i)$ denote the i -th configuration of a path σ , σ_i a suffix path starting with $\sigma(i)$, and Σ the set of paths. An execution is a path σ in Σ such that $\sigma(0) \in \mathcal{C}^0$.

2.2 Temporal Patterns

In this section, we also briefly recall the FTPL⁵ logic patterns introduced in [16]. In addition to configuration properties (*cp*) in *CP* mentioned above, the proposed logic contains external events (*ext*), as well as events from reconfiguration operations, temporal properties (*tpp*) together with trace properties (*trp*) embedded into temporal properties. Let $Prop_{FTPL}$ denote the set of the FTPL formulae obeying the FTPL grammar below. The FTPL semantics from [3] is summarized in Appendix A.

$$\begin{aligned}
\langle FTPL \rangle &::= \langle tpp \rangle \mid \langle events \rangle \mid cp \\
\langle tpp \rangle &::= \textbf{after} \langle events \rangle \langle tpp \rangle \mid \textbf{before} \langle events \rangle \langle trp \rangle \mid \langle trp \rangle \textbf{until} \langle events \rangle \mid \langle trp \rangle \\
\langle trp \rangle &::= \textbf{always} cp \mid \textbf{eventually} cp \mid \langle trp \rangle \wedge \langle trp \rangle \mid \langle trp \rangle \vee \langle trp \rangle \\
\langle events \rangle &::= \langle event \rangle, \langle events \rangle \mid \langle event \rangle \\
\langle event \rangle &::= ope \textbf{normal} \mid ope \textbf{exceptional} \mid ope \textbf{terminates} \mid ext
\end{aligned}$$

In the rest of the paper, let AE be the set of atomic events composed of atomic propositions from *CP* and of basic FTPL events. An *event* θ is an element of $\Theta = 2^{AE}$. Let us suppose that each component C_i of the component-based system has a local monitor M_i attached to it, from the set $\mathcal{M} = \{M_0, \dots, M_{n-1}\}$ of monitors⁶. Let us introduce the projection function Π_i to restrict events to the local view of the monitor M_i . For atomic events, $\Pi_i : 2^{AE} \rightarrow 2^{AE}$, and we write $AE_i = \Pi_i(AE)$. We assume $\forall i, j \leq n. i \neq j \Rightarrow AE_i \cap AE_j = \emptyset$ ⁷. Similarly, for events, we define $\Pi_i : 2^\Theta \rightarrow 2^\Theta$, with $\Theta_i = \Pi_i(\Theta)$, and we assume $\forall i, j \leq n. i \neq j \Rightarrow \Theta_i \cap \Theta_j = \emptyset$.

Let $ev : \mathcal{C} \rightarrow \Theta$ be a function to associate events with configurations. Given a configuration $\sigma(j)$ of a path σ with $j \geq 0$, the corresponding event is $\theta(j) = ev(\sigma(j))$. In this setting, an individual *behaviour* of a component C_i can be defined as a (finite or infinite) sequence of events $\theta_i = \theta_i(0) \cdot \theta_i(1) \cdots \theta_i(j) \cdots$ s.t.

⁵ FTPL stands for TPL (Temporal Pattern Language) prefixed by ‘F’ to denote its relation to Fractal-like components and to first-order integrity constraints over them.

⁶ Implemented as controllers in CSP||B, Fractal, FraSCaTi, etc.

⁷ For relations involving two components (like *Delegate* or *Parent*) we consider that only the parent component is *aware* of the relation. For the *Binding* relation, only the component owning the *required* (or *client*) interface is *aware* of the binding.

$\forall j \geq 0. \theta_i(j) = \Pi_i(ev(\sigma(j))) \in \Theta_i$, also called a *trace*. Finite (resp. infinite) traces over Θ are elements of Θ^* (resp. Θ^ω); the set of all traces is $\Theta^\infty = \Theta^* \cup \Theta^\omega$.

3 FTPL Progression and Urgency

This section provides the underpinnings to allow a decentralised evaluation of FTPL formulae. Inspired by definitions in [4], our notions of progression and urgency are adapted to the FTPL semantics: they take into account external and internal events as well as scopes of linear temporal patterns.

For decentralised evaluation of the FTPL formulae, instead of the set \mathbb{B}_4 as in [3], let us consider the set $\mathbb{B}_5 = \{\perp, \perp^p, \#, \top^p, \top\}$, where \perp, \top stand resp. for *false* and *true* values, \perp^p, \top^p for *potential false* and *potential true* values, and $\#$ for *unknown* value. We consider \mathbb{B}_5 together with the truth non-strict ordering relation \sqsubseteq_5 satisfying $\perp \sqsubseteq_5 \perp^p \sqsubseteq_5 \top^p \sqsubseteq_5 \top \sqsubseteq_5 \#$. On \mathbb{B}_5 we define two binary symmetric operations \sqcap_5, \sqcup_5 resp. as the minimum and maximum interpreted wrt. \sqsubseteq_5 . Thus, $(\mathbb{B}_5, \sqsubseteq_5)$ is a finite lattice but not a Boolean nor a de Morgan lattice. Let $\forall \varphi \in Prop_{FTPL}. \varphi \sqcap_5 \# = \varphi$. We write \sqcup and \sqcap instead of \sqcup_5 and \sqcap_5 when it is clear from the context. For any formula $\varphi \in Prop_{FTPL}$, let $\hat{\varphi}$ denote the value of φ in \mathbb{B}_5 .

In the context of a decentralised evaluation, each monitor may not be aware of information related to a given property and may be not able to evaluate it. This property is then written as a formula in terms of the current configuration. However, after the transition to the next configuration, such a formula may be not relevant. To compensate for this, we define the progression function to rewrite FTPL formulae in a way relevant to the next configuration of a path. Intuitively, given an FTPL formula and a set of atomic events, the progression function provides either the value of the property, if available, or the rewritten formula otherwise.

Definition 4 (Progression Function for *Events*). Let $\varepsilon, \varepsilon_1, \varepsilon_2 \in AE$, $e = e_1, e_2 \dots e_m$, a list of FTPL events from AE , and $\theta(i)$ an event. The progression function $P : Prop_{FTPL} \times \Theta \rightarrow Prop_{FTPL}$ is inductively defined by:

$$\begin{aligned} P(\varepsilon, \theta(i)) &= \top \text{ if } \varepsilon \in \theta(i), \perp \text{ otherwise} & ; & & P(\perp, \theta(i)) &= \perp \\ P(\varepsilon_1 \vee \varepsilon_2, \theta(i)) &= P(\varepsilon_1, \theta(i)) \vee P(\varepsilon_2, \theta(i)) & ; & & P(\perp^p, \theta(i)) &= \perp^p \\ P(\neg \varepsilon, \theta(i)) &= \neg P(\varepsilon, \theta(i)) & ; & & P(\top^p, \theta(i)) &= \top^p \\ P(e, \theta(i)) &= \bigvee_{1 \leq j \leq m} P(e_j, \theta(i)) & ; & & P(\top, \theta(i)) &= \top \end{aligned}$$

Let us now introduce, in order to establish progression formulae, the $\overline{\mathbf{X}}$ -operator that precedes an FTPL property to denote its evaluation at the configuration preceding the current one, i.e., $P(\overline{\mathbf{X}}\xi, \theta(i)) = P(\xi, \theta(i-1))$. We write $\overline{\mathbf{X}}^m \xi$ to denote $\overbrace{\overline{\mathbf{X}} \overline{\mathbf{X}} \dots \overline{\mathbf{X}}}^m \xi$. Also, when $m = 0$, $\overline{\mathbf{X}}^m \xi = \xi$.

Because of lack of room, the progression function is not given for every type of FTPL property. Instead, we provide a definition for the **always** trace property (Def. 5), lists of events (Def. 6), and the **before** temporal property (Def. 7). The reader can extrapolate these definitions for the remaining FTPL properties, using

the FTPL progressive semantics introduced in [3]. A whole definition for FTPL properties is given in Appendix B.

Definition 5 (Progression of the *always* FTPL trace property's evaluation formulae on a (suffix) path). Let cp be a configuration property and ϕ a trace property of the form $\phi = \mathbf{always} \ cp$. The progression function P for the **always** property on a (suffix) path is defined by:

$$P(\phi_{\sigma_k}, \theta(i)) = \begin{cases} P(cp, \theta(i)) \sqcap \top^p & \text{for } i = k \\ P(cp, \theta(i)) \sqcap P(\overline{\mathbf{X}}\phi_{\sigma_k}, \theta(i)) & \text{for } i > k \end{cases} \quad (1)$$

Definition 6 (Progression of FTPL list of events properties' evaluation formulae on a (suffix) path). Let e be a list of FTPL events, the progression function P for FTPL lists of events on a (suffix) path is defined by:

$$P(e_{\sigma_k}, \theta(i)) = \begin{cases} P(e, \theta(i)) & \text{for } i = k \\ P(e, \theta(i)) \sqcup (\top^p \sqcap P(\overline{\mathbf{X}}e_{\sigma_k}, \theta(i))) & \text{for } i > k \end{cases} \quad (2)$$

Definition 7 (Progression of the *before* FTPL temporal property's evaluation formulae on a (suffix) path). Let e be a list of FTPL events, trp a trace property, and β a temporal property of the form $\beta = \mathbf{before} \ e \ trp$. The progression function P for the **before** property on a (suffix) path is defined by:

$$P(\beta_{\sigma_k}, \theta(i)) = \begin{cases} \top^p & \text{for } i = k \\ F_{\mathcal{B}}(P(e_{\sigma_k}, \theta(i)), P(\overline{\mathbf{X}}trp_{\sigma_k}, \theta(i)), P(\overline{\mathbf{X}}\beta_{\sigma_k}, \theta(i))) & \text{for } i > k \end{cases} \quad (3)$$

where $F_{\mathcal{B}}$ is based on the FTPL progressive semantics and defined as follows:

$$F_{\mathcal{B}}(\varepsilon, trp, tpp) = \begin{cases} \top^p & \text{if } \varepsilon = \perp \\ \perp & \text{if } \varepsilon = \top \wedge trp \in \{\perp, \perp^p\} \\ tpp & \text{otherwise} \end{cases} \quad (4)$$

Example 1. Let be $\varphi = \mathbf{before} \ e \ trp$ where e is an FTPL list of events and trp a trace property. To evaluate φ at the configuration of index $i > 0$ on the suffix path σ_0 , let us set $P(e_{\sigma_0}, \theta(i)) = e_{\sigma_0}(i) = \top$ and $P(trp_{\sigma_0}, \theta(i-1)) = trp_{\sigma_0}(i-1) = \perp^p$. Then by Equalities (12) and (9) we have:

$$\begin{aligned} P(\varphi_{\sigma_0}, \theta(i)) &= F_{\mathcal{B}}(P(e_{\sigma_0}, \theta(i)), P(\overline{\mathbf{X}}trp_{\sigma_0}, \theta(i)), P(\overline{\mathbf{X}}\varphi_{\sigma_0}, \theta(i))) \\ &= F_{\mathcal{B}}(P(e_{\sigma_0}, \theta(i)), P(trp_{\sigma_0}, \theta(i-1)), P(\varphi_{\sigma_0}, \theta(i-1))) \\ &= F_{\mathcal{B}}(\top, \perp^p, P(\varphi_{\sigma_0}, \theta(i-1))) \\ &= \perp \end{aligned}$$

In order to perform evaluation in a decentralised manner, we define below the *Normalised Progression Form* (NPF) to describe the point up to which a formula should be developed, using the progression function.

Definition 8 (NPF). Let φ be an FTPL property and θ an event. A formula $P(\varphi, \theta)$ is in NPF if the $\overline{\mathbf{X}}$ -operator only precedes atomic events.

Theorem 1 (Existence of NPF). *Let φ be an FTPL property and θ an event. Every $P(\varphi, \theta)$ can be rewritten into an equivalent⁸ formula in NPF.*

Proof. The proof is by induction on the indexes of the events (i.e., on the trace) using Definitions 5 to 7 (and definitions for the remaining FTPL properties).

Example 2. Let be $\varphi = \mathbf{before} \ e \ trp$, $e = a, b$, and $trp = \mathbf{always} \ cp$, where a and b are FTPL events s.t. a, b , and $cp \in CP$ are atomic events. The resulting formula in NPF is obtained using Equation 12.

$$\begin{aligned}
P(\varphi_{\sigma_0}, \theta(0)) &= \top^p \\
P(\varphi_{\sigma_0}, \theta(1)) &= F_B(P(e_{\sigma_0}, \theta(1)), P(\overline{\mathbf{X}}trp_{\sigma_0}, \theta(1)), P(\overline{\mathbf{X}}\varphi_{\sigma_0}, \theta(1))) \\
&= F_B(P(e, \theta(1)) \sqcup (\top^p \sqcap P(\overline{\mathbf{X}}e_{\sigma_0}, \theta(1))), P(trp_{\sigma_0}, \theta(0)), P(\varphi_{\sigma_0}, \theta(0))) \\
&= F_B(P(a, \theta(1)) \sqcup P(b, \theta(1)) \sqcup (\top^p \sqcap P(e_{\sigma_0}, \theta(0))), P(cp, \theta(0)) \sqcap \top^p, \top^p) \\
&= F_B(P(a, \theta(1)) \sqcup P(b, \theta(1)) \sqcup (\top^p \sqcap P(e, \theta(0))), P(cp, \theta(0)) \sqcap \top^p, \top^p) \\
&= F_B(P(a, \theta(1)) \sqcup P(b, \theta(1)) \sqcup (\top^p \sqcap (P(a, \theta(0)) \sqcup P(b, \theta(0)))), P(cp, \theta(0)) \sqcap \top^p, \top^p) \\
&= F_B(P(a, \theta(1)) \sqcup P(b, \theta(1)) \sqcup (\top^p \sqcap (P(\overline{\mathbf{X}}a, \theta(1)) \sqcup P(\overline{\mathbf{X}}b, \theta(1)))), P(\overline{\mathbf{X}}cp, \theta(1)) \sqcap \top^p, \top^p)
\end{aligned}$$

As in [4] for LTL, a monitor M_j for the component C_j accepts as input an event $\theta(i)$ and FTPL properties. Applying Def. 4 to atomic events could lead to wrong results in a decentralised context. For example, if $\varepsilon \notin \theta(i)$ holds locally for the monitor M_j it could be due to the fact that $\varepsilon \notin AE_j$. The decentralised progression rule should be adapted by taking into account a local set of atomic events. Hence, the progression rule for atomic events preceded by the $\overline{\mathbf{X}}$ -operator is given below.

$$P(\overline{\mathbf{X}}^m \varsigma, \theta(i), AE_j) = \begin{cases} \top & \text{if } \varsigma = \varsigma' \text{ for some } \varsigma' \in AE_j \cap \Pi_j(\theta(i-m)), \\ \perp & \text{if } \varsigma = \varsigma' \text{ for some } \varsigma' \in AE_j \setminus \Pi_j(\theta(i-m)), \\ \overline{\mathbf{X}}^{m+1} \varsigma & \text{otherwise.} \end{cases} \quad (5)$$

We complete the specification of the progression function with the special symbol $\# \notin AE$ for which the progression is defined by $\forall j. P(\#, \theta, AE_j) = \#$. Finally, among different formulae to be evaluated, the notion of urgency allows determining a set of urgent formulae. In a nutshell, the urgency of a formula in NPF is 0 if the formula does not contain any $\overline{\mathbf{X}}$ -operator or the value of the greatest exponent of the $\overline{\mathbf{X}}$ -operator. Using formulae in NPF, any sub-formula ς following an $\overline{\mathbf{X}}$ -operator is atomic ($\exists j. \varsigma \in AE_j$) and can only be evaluated by a single monitor M_j . A formal definition of urgency is as follows.

Definition 9 (Urgency). *Let φ be a FTPL formula, and $\Upsilon : Prop_{FTPL} \rightarrow \mathbb{N}^{\geq 0}$ an inductively defined function assigning a level of urgency to a FTPL formula as follows:*

$$\begin{aligned}
\Upsilon(\varphi) = \text{match } \varphi \text{ with } & \varphi_1 \sqcup \varphi_2 \mid \varphi_1 \sqcap \varphi_2 \mid F_A(_, \varphi_1, \varphi_2) & \rightarrow \max(\Upsilon(\varphi_1), \Upsilon(\varphi_2)) \\
& \mid F_B(\varphi_1, \varphi_2, \varphi_3) & \rightarrow \max\{\Upsilon(\varphi_i) \mid 1 \leq i \leq 3\} \\
& \mid F_U((\varphi_1, \varphi_2), (\varphi_3, \varphi_4), \varphi_5) & \rightarrow \max\{\Upsilon(\varphi_i) \mid 1 \leq i \leq 5\} \\
& \mid \overline{\mathbf{X}}\varphi & \rightarrow 1 + \Upsilon(\varphi) \\
& \mid _ & \rightarrow 0
\end{aligned}$$

⁸ wrt. the semantics.

where F_A (resp. F_U) denotes the FTPL temporal property **after** (resp. **until**), and F_B stands for the **before** property in Def. 7.

A formula φ is said to be more urgent than a formula ψ if and only if $Y(\varphi) > Y(\psi)$ holds. A formula φ where $Y(\varphi) = 0$ holds is said to be not urgent. Furthermore, we define the set of urgent sub-formulae of φ by $sus(\varphi)$, with $sus : Prop_{FTPL} \rightarrow 2^{Prop_{FTPL}}$ defined inductively as follows:

$$\begin{aligned}
 sus(\varphi) = \text{match } \varphi \text{ with } & \varphi_1 \sqcup \varphi_2 \mid \varphi_1 \sqcap \varphi_2 \mid F_A(_, \varphi_1, \varphi_2) & \rightarrow sus(\varphi_1) \cup sus(\varphi_2) \\
 & \mid F_B(\varphi_1, \varphi_2, \varphi_3) & \rightarrow \cup_{1 \leq i \leq 3} sus(\varphi_i) \\
 & \mid F_U((\varphi_1, \varphi_2), (\varphi_3, \varphi_4), \varphi_5) & \rightarrow \cup_{1 \leq i \leq 5} sus(\varphi_i) \\
 & \mid \neg \varphi & \rightarrow sus(\varphi) \\
 & \mid \overline{\mathbf{X}}\varphi & \rightarrow \{\overline{\mathbf{X}}\varphi\} \\
 & \mid _ & \rightarrow \emptyset
 \end{aligned}$$

4 Decentralised Evaluation Problem

As FTPL patterns contain conjunctions or disjunctions of properties over components' parameters and relations, the evaluation of temporal patterns in a decentralised manner makes sense. Section 4.1 addresses the temporal pattern decentralised evaluation problem on a reconfiguration path by providing an algorithm for such an evaluation in \mathbb{B}_4 . Its properties are studied in Section 4.2.

4.1 Problem Statement and Local Monitor Algorithm

Let $\hat{\varphi}_{\sigma_k}(s)$ denote the value of φ at configuration of index s on the suffix path σ_k . While considering components with their monitors, because of a decentralised fashion, the evaluation of $\varphi_{\sigma_k}(s)$ by a monitor M_i may be delayed to configuration $\sigma(t)$ with $t > s$, and the progression comes into play. In this case, let ${}_i\hat{\varphi}_{\sigma_k}^s(t)$ denote the decentralised formula as progressed to the configuration $\sigma(t)$ by M_i , for the evaluation of φ started at configuration $\sigma(s)$. Therefore, we consider the following decision problem.

Temporal Pattern Decentralised Evaluation on a Path (TPDEP).

Input: an FTPL temporal property φ , a suffix path σ_k with $k \geq 0$, a configuration $\sigma(s)$ with $s \geq k$, and a number $n = |\mathcal{M}|$ of monitors.

Output: $i, j < n$, and ${}_i\hat{\varphi}_{\sigma_k}^s(s+j) \in \mathbb{B}_4$, the value of φ at $\sigma(s+j)$ by M_i .

We consider as the basic TPDEP case the situation when only *run* operations occur after the TPDEP problem input, and until an output is returned, communications between monitors being covered by *run* operations.

The idea of a decentralised evaluation is as follows. Similarly to [4], at the configuration $\sigma(t)$, if ${}_i\varphi^s(t)$ cannot be evaluated in \mathbb{B}_4 , a monitor M_i progresses its current formula ${}_i\varphi^s(t)$ to ${}_i\varphi^s(t+1) = P({}_i\varphi^s(t), \theta(t), AE_i)$ and sends it to a monitor that can evaluate its most urgent sub-formula. After ${}_i\varphi^s(t+1)$ is sent,

M_i sets ${}_i\varphi^s(t+1) = \#$. When M_i receives one or more formulae from others monitors, each of them is added to the current formula using the \sqcap operator.

Unlike [4], where LTL decentralised monitoring determines the *steady* value of a property in \mathbb{B}_2 , our decentralised method allows values of FTPL properties in \mathbb{B}_4 to vary at different configurations, depending notably on the property scopes and on external events. To this end, a result in \mathbb{B}_4 obtained by a monitor is broadcast to other monitors, allowing them to maintain a complete but bounded history that can be used to invoke the TPDEP problem at the following configurations.

To answer the TPDEP problem, we propose the **LDMon** algorithm displayed in Fig. 2. It takes as input the index i of the current monitor, its set AE_i of atomic events, the index s of the current configuration, an FTPL temporal property φ to be evaluated, and the index k of the suffix path on which φ is supposed to be evaluated. An integer variable t indicates the index of the current configuration as it evolves. The algorithm broadcasts to all monitors, as soon as it is determined, the result of the evaluation of φ in \mathbb{B}_4 . We chose this method to transmit the results because we prefer to focus on the feasibility of a decentralised evaluation of temporal patterns and we consider that the transmission of result is a related issue outside of the scope of this paper.

Three functions are used in this algorithm: *a*) **send**(φ), sends φ (as well as its sub-formulae evaluated at the current configuration) to monitor M_j (different from the current monitor) where ψ is the most urgent sub-formula⁹ such that $Prop(\psi) \subseteq AE_j$ holds, with $Prop : Prop_{FTPL} \rightarrow 2^{AE}$ yielding the set of events of an FTPL formula; *b*) **receive**($\{\varphi, \dots\}$), receives formulae sent (and broadcast) by other monitors; and *c*) **broadcast**(φ), broadcasts φ to all other monitors.

As long as an evaluation of φ in \mathbb{B}_4 is not obtained (line 11), the **LDMon** algorithm loops in the following way: the evaluation formula is progressed to the next configuration (line 12) and the current configuration index t is incremented (line 13). If at least one event of the current formula belongs to the set of atomic events AE_i ($Prop(\varphi) \cap AE_i \neq \emptyset$) and if no progressed formula was sent (or if such a formula was sent and at least one from another monitor was received) at the previous configuration (${}_i\varphi_{\sigma_k}^s(t) \neq \#$), the progressed formula is sent to the monitor that can solve its most urgent sub-formula (line 15) and is set to $\#$ (line 16). Progressed formulae (and broadcast results) from other monitors are received (line 18) and are combined to the local formula using the \sqcap -operator (line 19). If the result is not in \mathbb{B}_4 , the loop continues, otherwise if the result of the formula has not already been provided by another monitor (line 21), the result is broadcast (line 22) and returned (line 24).

4.2 Correctness, Uniqueness, and Termination

In this section several properties of the **LDMon** algorithm are studied. Proposition 1, below, guarantees that the **LDMon** algorithm provides an output within a finite number of configurations, communications being covered by *run* operations.

⁹ In the case where there are two or more equally urgent formulae, φ is sent to a monitor determined by an arbitrary order with the function $Mon : \mathcal{M} \times 2^{AE} \rightarrow \mathcal{M}$. $Mon(M_i, AE') = M_{j_{min}}$ s.t. $j_{min} = \min(j \in [1, n] \setminus \{i\} | AE' \cap AE_j \neq \emptyset)$.

```

1  (*LDMon*)
2  Input
3     $i$       (*current monitor index*)
4     $AE_i$    (*current monitor atomic events*)
5     $s$       (*input configuration index*)
6     $\varphi$      (*FTPL property to evaluate*)
7     $\sigma_k$  (*suffix path*)
8  Variables
9     $t := s : integer$ 
10 Begin
11 | WHILE ( ${}_i\hat{\varphi}_{\sigma_k}^s(t) \notin \mathbb{B}_4$ ) DO
12 | |  ${}_i\varphi_{\sigma_k}^s(t+1) := P({}_i\varphi_{\sigma_k}^s(t), \theta_i(t), AE_i)$ 
13 | |  $t := t + 1$ 
14 | | IF ( $Prop(\varphi) \cap AE_i \neq \emptyset \wedge {}_i\varphi_{\sigma_k}^s(t) \neq \#$ ) DO
15 | | | send( ${}_i\varphi_{\sigma_k}^s(t)$ )
16 | | |  ${}_i\varphi_{\sigma_k}^s(t) = \#$ 
17 | | FI
18 | | receive( $\{{}_j\varphi_{\sigma_k}^s(t)\}_{j \neq i}$ )
19 | |  ${}_i\varphi_{\sigma_k}^s(t) := {}_i\varphi_{\sigma_k}^s(t) \sqcap \prod_{j \neq i} {}_j\varphi_{\sigma_k}^s(t)$ 
20 | ENDWHILE
21 | IF ( $\forall j \neq i. {}_j\hat{\varphi}_{\sigma_k}^s(t) \notin \mathbb{B}_4$ ) DO
22 | | broadcast( ${}_i\varphi_{\sigma_k}^s(t)$ )
23 | FI
24 | RETURN  ${}_i\hat{\varphi}_{\sigma_k}^s(t)$ 
25 End

```

Fig. 2: Algorithm LDMon

Proposition 1 (Existence). *Let $\varphi \in Prop_{FTPL}$, σ_k a suffix path, $k \geq 0$. For a given configuration $\sigma(s)$ with $s \geq k$, when using a number $n = |\mathcal{M}|$ of monitors, the LDMon algorithm provides an output such that $\exists i, j. i, j < n \wedge {}_i\hat{\varphi}_{\sigma_k}^s(s+j) \in \mathbb{B}_4$.*

Proof. (Sketch.) Let M_0, M_1, \dots, M_{n-1} be n monitors. At a given configuration of index s , if one of the monitors $M_i \in \mathcal{M}$ is able to evaluate its formula in \mathbb{B}_4 , the proposition holds with $j = 0$. Otherwise, each monitor M_i ($0 \leq i < n$) progresses its formula ${}_i\varphi_{\sigma_k}^s(s)$ into ${}_i\varphi_{\sigma_k}^s(s+1)$ and sends it to another monitor, according to *Mon*, able to answer its most urgent sub-formula.

We assume that ${}_i\varphi_{\sigma_k}^s(s+1)$ is sent to the monitor $M_{i_2 \neq i_1}$. At the next configuration of index $s+1$, the monitor M_{i_2} receives ${}_i\varphi_{\sigma_k}^s(s+1)$ and combines it with ${}_{i_2}\varphi_{\sigma_k}^s(s+1)$ as well as other formulae (if any) received from other monitors using the \sqcap -operator. If one of these formulae (or a sufficient number of sub-formulae) can be evaluated in \mathbb{B}_4 , the proposition holds with $j = 1$ and ${}_i\hat{\varphi}_{\sigma_k}^s(s+1)$. Otherwise, each monitor M_i progresses the formula ${}_i\varphi_{\sigma_k}^s(s+1)$ into ${}_i\varphi_{\sigma_k}^s(s+2)$ and sends it to another monitor according to *Mon* which is able to answer its most urgent sub-formula.

We assume that ${}_{i_2}\varphi_{\sigma_k}^s(s+2)$ is sent to the monitor M_{i_3} with $i_3 \neq i_2$. Also $i_3 \neq i_1$ because previously, all sub-formulae of ${}_{i_1}\varphi_{\sigma_k}^s(s+1)$ that could be solved using the set of atomic events AE_{i_1} were already solved. This way, the problem is reduced from n to $n-1$ monitors. Since for a single monitor the output of the algorithm is $\hat{\varphi}_{\sigma_k}(s)$ with $j=0$, we can infer that for n monitors, there is at least one monitor M_{i_0} such that ${}_{i_0}\hat{\varphi}_{\sigma_k}^s(s+j) \in \mathbb{B}_4$ with $j < n$. \square

As explained before, when evaluating $\varphi_{\sigma_k}(s)$, the formula ${}_{i}\varphi_{\sigma_k}^s(t)$ at configuration of index t by M_i either has a result ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) \in \mathbb{B}_4$ or progresses to $\#$. The latter is written ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) = \#$. Thus ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) \in \mathbb{B}_5$.

Theorem 2 (Semantic Correctness). ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) \neq \# \Leftrightarrow {}_{i}\hat{\varphi}_{\sigma_k}^s(t) = \hat{\varphi}_{\sigma_k}(s)$.

Proof. (Sketch.)

\Rightarrow If ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) \neq \#$, a result has been obtained in \mathbb{B}_4 , otherwise ${}_{i}\hat{\varphi}_{\sigma_k}^s(t)$ would equal $\#$. Therefore, we only have to verify that the progression function of Def. 4 to 7 (and definitions for the remaining FTPL properties) matches the FTPL semantics in \mathbb{B}_4 as defined in [3]. It is done by induction on the path length.

\Leftarrow ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) = \hat{\varphi}_{\sigma_k}(s) \Rightarrow {}_{i}\hat{\varphi}_{\sigma_k}^s(t) \in \mathbb{B}_4 \Rightarrow {}_{i}\hat{\varphi}_{\sigma_k}^s(t) \neq \#$.

Corollary 1 (Uniqueness). *If ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) \neq \#$ and ${}_{j}\hat{\varphi}_{\sigma_k}^s(t) \neq \#$ for $i \neq j$, then ${}_{i}\hat{\varphi}_{\sigma_k}^s(t) = {}_{j}\hat{\varphi}_{\sigma_k}^s(t)$.*

Corollary 2 (Generalised Uniqueness). *Let be ${}_S\varphi_{\sigma_k}^s(t) = \bigcap_{i \in S} {}_{i}\varphi_{\sigma_k}^s(t)$ for $S \subseteq [1, n]$. If ${}_S\hat{\varphi}_{\sigma_k}^s(t) \neq \#$ then for all $j \in S$, ${}_j\hat{\varphi}_{\sigma_k}^s(t) \neq \#$ implies ${}_j\hat{\varphi}_{\sigma_k}^s(t) = {}_S\hat{\varphi}_{\sigma_k}^s(t)$.*

Corollary 2 allows a monitor to simplify the combination of formulae with the operator \sqcap . For a given property, a conjunction in \mathbb{B}_4 of formulae received from other monitors with the formula of the current monitor can be replaced by any of these formulae provided that its value is different from $\#$.

Example 3. Let us consider again $\varphi = \text{before } e \text{ trp}$. Let A, B , and C be the components with their respective monitors M_A, M_B , and M_C such that ${}_{i}\varphi^s(t) = F_{\mathcal{B}}({}_{i}e^s(t), {}_{i}\text{trp}^{s-1}(t), {}_{i}\varphi^{s-1}(t))$ for $i \in \{A, B, C\}$ (Def. 7). Let us assume $\varphi(s) = F_{\mathcal{B}}(e(s), \text{trp}(s-1), \varphi(s-1))$, with $\varphi(s)$, $e(s)$, $\text{trp}(s-1)$, and $\varphi(s-1)$ being evaluated in \mathbb{B}_4 . By Corollary 2, $e(s) = {}_Ae^s(t) \sqcap {}_Be^s(t) \sqcap {}_Ce^s(t)$ (resp. $\text{trp}(s-1) = {}_A\text{trp}^{s-1}(t) \sqcap {}_B\text{trp}^{s-1}(t) \sqcap {}_C\text{trp}^{s-1}(t)$, $\varphi(s-1) = {}_A\varphi^{s-1}(t) \sqcap {}_B\varphi^{s-1}(t) \sqcap {}_C\varphi^{s-1}(t)$) if it exists at least one i such that the value of ${}_{i}e^t(s)$ (resp. ${}_{i}\text{trp}^{s-1}(t)$, ${}_{i}\varphi^{s-1}(t)$) is in \mathbb{B}_4 ; in this case, ${}_ie^s(t) = e(s)$ (resp. ${}_{i}\text{trp}^{s-1}(t) = \text{trp}(s-1)$, ${}_{i}\varphi^{s-1}(t) = \varphi(s-1)$).

For example, if ${}_A\varphi^s(t) = F_{\mathcal{B}}(\top, \phi, {}_A\varphi^{s-1}(t))$, ${}_B\varphi^s(t) = F_{\mathcal{B}}(\varepsilon, \top^p, {}_B\varphi^{s-1}(t))$, and ${}_C\varphi^s(t) = F_{\mathcal{B}}(\epsilon, \psi, \top^p)$, with ϕ , ${}_A\varphi^{s-1}(t)$, ε , ${}_B\varphi^{s-1}(t)$, ϵ , and ψ not being evaluated in \mathbb{B}_4 . It implies ${}_Ae^s(t) = e(s) = \top$ (resp. ${}_B\text{trp}^{s-1}(t) = \text{trp}(s-1) = \top^p$, ${}_C\varphi^{s-1}(t) = \varphi(s-1) = \top^p$) and $\varphi(s) = F_{\mathcal{B}}(\top, \top^p, \top^p) = \top^p$.

Proposition 2 (Correctness and Uniqueness). *The output provided by the LDMon algorithm answers the TPDEP problem. For a given configuration $\sigma(s)$, this answer is unique.*

Proof. (Sketch.) By Proposition 1 LDMon provides an output ${}_i\hat{\varphi}_{\sigma_k}^s(s+j)$ for at least one monitor M_i , within a finite number j of configurations. By Theorem 2 this output answers the TPDEP problem. Furthermore Corollary 1 establishes that for any i_0 , if ${}_{i_0}\hat{\varphi}_{\sigma_k}^s(s+j)$ is the output of the LDMon algorithm for the monitor M_{i_0} then ${}_{i_0}\hat{\varphi}_{\sigma_k}^s(s+j) = {}_i\hat{\varphi}_{\sigma_k}^s(s+j)$. \square

Proposition 3 (Termination). *The LDMon algorithm always terminates, either at the configuration when an output is provided or at the next one. Furthermore, the number of configurations needed to reach a result is at most $|\mathcal{M}|$.*

Proof. (sketch) Propositions 1 and 2 establish that the LDMon algorithm terminates and answers the TPDEP problem for at least one monitor M_i after a finite number of reconfigurations $j < |\mathcal{M}|$. Such monitor M_i broadcasts the result to all other monitors before finishing (line 22 of the LDMon algorithm, Fig. 2). This enables any monitor for which the LDMon algorithm did not finish at configuration $s+j$ to receive the result of the broadcast and to finish its instance of the LDMon algorithm at configuration $s+j+1 \leq s+|\mathcal{M}|$. \square

In general, decentralised algorithms tend to be very hard for creating a consensus and moreover they require significant communication overhead. Let us emphasize the fact that Proposition 2 guarantees the correctness and uniqueness of a result, which implies such a consensus. As a consequence of Propositions 2 and 3 adaptation policies relying on the decentralised evaluation of FTPL temporal properties can be applied to component-based systems for their dynamic reconfiguration at runtime.

Let us now discuss communication overhead. We consider a component-based system of N components reporting their status in \mathbb{B}_4 to a central controller at each configuration as described for in [3]. In the centralised context, thanks to the progressive semantics, the evaluation of a given FTPL property φ would mean that N messages should be sent to conclude in \mathbb{B}_4 . With the decentralised approach, assuming that atomic events of φ would be distributed among n components ($n \leq N$), we would need, at most, $n^2 - 1$ messages to evaluate φ .

This means that to evaluate a formula involving $n = 10$ components of a component-based system of $N = 100$ components, in the worst case the decentralised fashion would need 99 messages versus 100 for the centralised approach to evaluate φ which is a ratio of 99%. If, however, the total number N of components of the system is much greater than the number n of components involved in the evaluation of φ , the communication overhead ratio can be even lower (e.g., 9.9% for $N = 1000$). Reciprocally, if a great proportion of the system is involved in the property to evaluate, the centralised method would lead to better results. Let q be such a proportion, i.e., $n = qN$, the communication overhead ratio is $Nq^2 - 1/N$.

This is different from the result in [4] where the decentralised algorithm outperforms its centralised counterpart by a proportion of 1 to 4 in terms of communication overhead, to conclude in \mathbb{B}_2 . Such a difference is due to the fact that in our case, as soon as a property is evaluated in \mathbb{B}_4 for a given configuration of the path, another evaluation is initiated for another configuration. Nevertheless, we have better results while monitoring only components concerned with the temporal property, that can be determined syntactically. To sum up, our approach is suitable for systems with a large number of components when the FTPL property to evaluate involves a small proportion of them.

5 Implementation and Experiment

This section describes how the **LDMon** algorithm has been implemented within the *GROOVE* graph transformation tool [5]. This implementation is then used to experiment with a case study.

5.1 Implementing with *GROOVE*

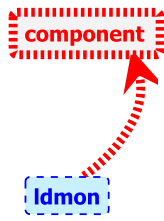


Fig. 3: Rule **removeOrphanMon**

GROOVE uses simple graphs for modelling the structure of object-oriented systems at design-time, compile-time, and run-time. Graphs are made of nodes and edges that can be labelled. Graph transformations provide a basis for model transformation or for operational semantics of systems. Graphs are transformed by rules consisting of *a*) patterns that must be present (resp. absent) for the rule to apply, *b*) elements (nodes and edges) to be added (resp. deleted) from the graph, and *c*) pairs of nodes to be merged. Colour and shape coding allow these rules to be easily represented. For example, our implementation uses the graph rule **removeOrphanMon** represented Fig. 3 that can be interpreted as follows:

a) The red (dashed fat) “embargo” elements, representing a node of type *component* and an edge defining a monitoring relation between monitors, of type *ldmon*, and *components*, must be absent, *b*) the blue (dashed thin) “eraser” element, representing a node of type *ldmon*, must be present, and *c*) if both conditions are satisfied, the blue (dashed thin) element is deleted. This means that if a monitor of type *ldmon* is not monitoring a *component*, the monitor node, *ldmon*, must be deleted. The reader interested in *GROOVE* is referred to [5].

Our implementation uses the *GROOVE* typed mode to guarantee that all graphs are well-typed. It consists of generic types and graph rules that can manage assigned priorities in such a way that a rule is applied only if no rule of higher priority matches the current graph. The input is a graph containing an FTPL formula and a component-based system, both represented using the model presented in Sec. 2. Figure 4 shows a screenshot of *GROOVE* displaying, in the main panel, a graph modelling the *location* component-based system used in the case study below. *Components* are represented in blue, *Required* (resp. *Provided*)

Interfaces in magenta (resp. red), *Parameters* in black, and both *ITypes* and *PTypes* in grey. The top left panel shows graph rules ordered by priority, whereas the bottom left panel contains *GROOVE* types.

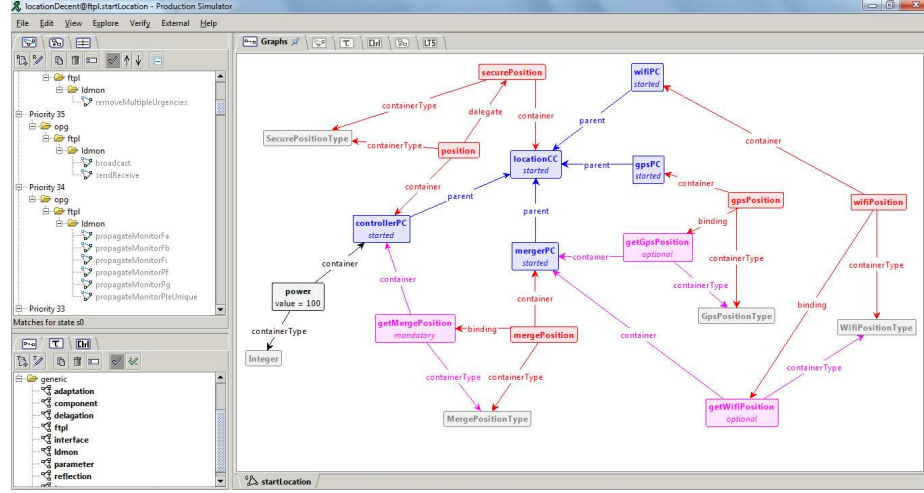


Fig. 4: Model of the location component-based system displayed with *GROOVE*

5.2 Case Study

In this section we illustrate the *LDMon* algorithm with an example of a location composite component, and afterwards we provide several details on its implementation in *GROOVE*. The location system is made up of different positioning systems, like GPS or Wi-Fi, a merger and a controller. Thanks to adaptation policies with temporal patterns, the location composite component can be modified to use either GPS or Wi-Fi positioning systems, depending on some properties, such as available energy, occurrences of indoor/outdoor positioning external events, etc. For example, when the level of energy is low, if the vehicle is in a tunnel where there is no GPS signal, it would be useful to remove the GPS component (cf. Fig. 5). To save energy, this component may not be added back before the level of energy reaches an acceptable value.

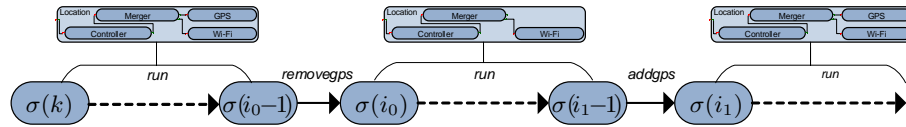


Fig. 5: Representation of the suffix configuration path σ_k

This example has been fully implemented with *GROOVE* together with adaptation policies. Let G be the *GROOVE* graph representing this example. Let us consider the FTPL temporal property $\varphi = \mathbf{after\ removegps\ normal\ (eventually\ (power \geq 33)\ before\ addgps\ normal)}$, which can be written as $\varphi = \mathbf{after\ } e_0\ \phi$, with $e_0 = \mathbf{removegps\ normal}$, $\phi = \mathbf{trp\ before\ } e_1$, $\mathbf{trp} = \mathbf{eventually\ } cp$, $e_1 = \mathbf{addgps\ normal}$, and $cp = (power \geq 33)$.

Intuitively, φ represents the requirement “After the GPS component has been removed, the level of energy has to be greater than 33% before this component is added back”. Figure 6 shows how φ is represented in our implementation.

Let M_c , M_m , M_g , and M_w be four monitors pertaining respectively to the controller, merger, GPS, and Wi-Fi components. Monitor M_c has access to the value of the configuration property **power_ge_33** (\top if $power \geq 33$, or \perp otherwise) while M_m is aware of the values of **addgps_normal** (resp. **removegps_normal**) which are \top at the configuration following the addition (resp. removal) of the GPS component, or \perp otherwise. Since monitors, M_g and M_w do not have access to any atomic event having an influence on the evaluation of φ (i.e., $Prop(\varphi) \cap AE_g = Prop(\varphi) \cap AE_w = \emptyset$), M_g and M_w do not send messages, which has a beneficial effect on the communication overhead. In our implementation, each monitor is a subgraph of G containing the monitored component via an edge named *monitor*. Communications between monitors are represented by edges named *sentreceived* and *broadcast*. Recall that in the model, communications between monitors are covered by *run* operations as they do not directly affect the system’s architecture. Let us consider a reconfiguration path σ representing the sequences of configurations of the location composite component where the transitions between configurations are reconfiguration operations. In the suffix path σ_k displayed in Fig. 5, we suppose that all the reconfiguration operations are *run*, except between $\sigma(i_0 - 1)$ and $\sigma(i_0)$ (resp. $\sigma(i_1 - 1)$ and $\sigma(i_1)$), where it is **removegps** (resp. **addgps**). During runtime, an adaptation controller—in charge of the application of adaptation policies—

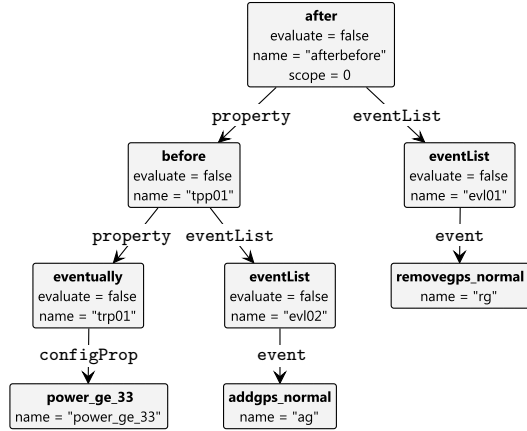


Fig. 6: Representation of the φ FTPL property

needs to evaluate FTPL properties. To do so, the adaptation controller uses the **LDMon** algorithm to evaluate φ . When a result is returned by a monitor, the most suitable adaptation policy is applied, and the algorithm is used again to evaluate φ at the following configuration, where it may vary because of the scope, for example. In the following, we describe how, at each configuration of index s

($k \leq s \leq i_1$) the adaptation controller requests the evaluation of φ to the monitors using the **LDMon** algorithm, and receives the answer at the same configuration $\sigma(s)$. In Fig. 7 green (dashed) arrows represent broadcast communications, blue (plain) arrows represent formulae being sent, and red (dotted) arrows indicate that *a*) the destination component is able to solve the most urgent sub-formula of the source component and *b*) no communication is made between these components. Because neither M_g nor M_w has access to atomic events impacting their formulae, they do not send any message during the run described below.

At configuration $\sigma(k)$, since M_m can evaluate $e_0 = \perp$, by definition of the **after** FTPL property, $m\hat{\varphi}_{\sigma_k}^k(k) = \top^p$ is established and broadcast. Other monitors progress their formulae and determine that the most urgent sub-formula can be solved by M_m (Fig. 7a); consequently, M_c sends its formulae¹⁰ to M_m .

At every configuration $\sigma(s)$ for $k+1 \leq s \leq i_0-1$, since e_0 does not occur, the decentralised evaluation consists in evaluating φ by M_m that returns and broadcasts the result. Other monitors receive the result from the previous configuration broadcast by M_m ¹¹. They also progress their current formulae, which cause M_c to send its formulae to M_m . This is displayed in Fig. 7b, where F_A represents the FTPL temporal property **after** in the same way F_B does for the **before** property in Def. 7.

At configuration $\sigma(i_0)$, the event $e_0 = \text{removegps normal}$, signifying the GPS normal removal, occurs. The M_m monitor, being aware of this occurrence, evaluates φ : $m\hat{\varphi}_{\sigma_k}^{i_0}(i_0) = m\hat{\phi}_{\sigma_{i_0}}^{i_0}(i_0) = \top^p$ because the “**before**” FTPL pattern is defined to be \top^p at the first configuration of its scope. The result is then returned and broadcast. In the meantime, M_c and M_w receive the result broadcast at the previous configuration and M_c sends its formulae to M_m .

At configuration $\sigma(s)$ for $i_0+1 \leq s \leq i_1-1$, because e_0 occurred once, M_m computes $m\hat{\varphi}_{\sigma_k}^s(s) = m\hat{\phi}_{\sigma_{i_0}}^s(s) = \top^p$, since $\phi = \text{trp before } e_1$ and e_1 has not

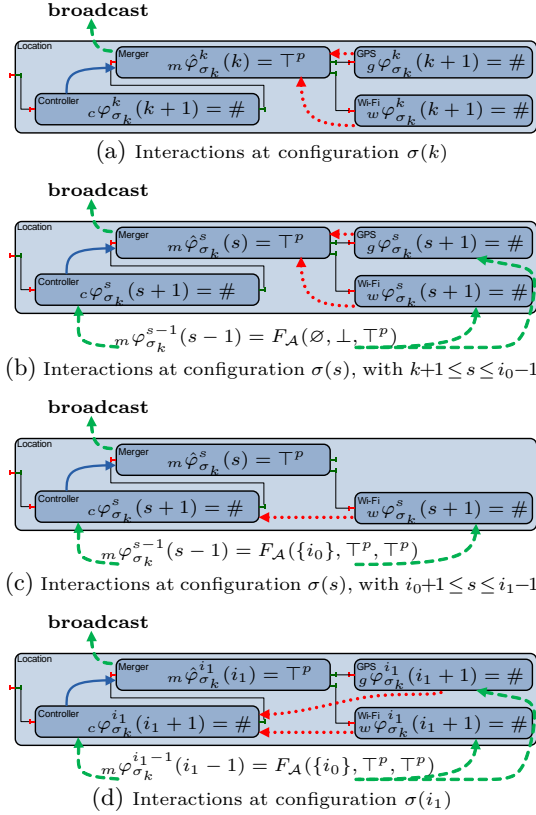


Fig. 7: Interactions between monitors

¹⁰ The formula to evaluate as well as its sub-formulae evaluated at the current state.

¹¹ This allows all monitors to keep a history of $|\mathcal{M}| + 1$ configurations.

yet occurred; the result is then returned and broadcast. M_c and M_w receive the result broadcast at the previous configuration which contains, as a sub-formula, the information that e_0 occurred at configuration $\sigma(i_0)$. The formula progressed by M_c contains ${}_c\phi_{\sigma_{i_0}}^s(s+1) = F_{\mathcal{B}}(\overline{\mathbf{X}}e_1, \hat{trp}_{\sigma_{i_0}}^{s-1}(s), \top^p)$. We suppose that there is a configuration $\sigma(s')$ s.t. $s' > i_0$, where the power rises over 33%, i.e., $cp = (\text{power} \geq 33) = \top$ and then $\hat{trp}_{\sigma_{i_0}}^s(s) = cp \sqcup \hat{trp}_{\sigma_{i_0}}^{s-1}(s) = \top$ for $s \geq s'$. In this case, the set of formulae M_c sends to M_m (Fig. 7c) contains $F_{\mathcal{B}}(\overline{\mathbf{X}}e_1, \top, \top^p)$ and $trp_{\sigma_{i_0}}^s(s)$.

At configuration $\sigma(i_1)$, $e_1 = \text{addgps terminates}$ just occurred. We assume that the reconfiguration terminated normally and that the GPS component was added. M_c , M_g , and M_w receive the result broadcast at the previous configuration. M_c and M_w behave in a way similar than above at configuration $\sigma(s)$ with $i_0 + 1 \leq s \leq i_1 - 1$, whereas M_g behaves like M_w . Finally, M_m evaluates its formula to ${}_m\hat{\phi}_{\sigma_k}^{i_1}(i_1) = {}_m\hat{\phi}_{\sigma_{i_0}}^{i_1}(i_1) = F_{\mathcal{B}}(\top, \overline{\mathbf{X}}trp_{\sigma_{i_0}}, \top^p) = \top^p$ using the fact that the sub-formula $\overline{\mathbf{X}}trp_{\sigma_{i_0}}$ was sent by M_c at the previous configuration. This result answers correctly the TPDEP.

Back to the implementation, Table 1 gives information on the *GROOVE* graph rules for the case study. The columns show, from left to right, the possible values of the index of the considered configurations, the number of graph rules used, the reconfiguration occurring (if any), and the part of the FTPL formula that must be evaluated to obtain a result in \mathbb{B}_4 . At configuration $\sigma(k)$ 85 rules are used, rules concerning the evaluation of FTPL events are the ones used the most; as long as the event **removegps normal** has not occurred yet, only the evaluation of the part “**after removegps normal**...” of the formula is needed to obtain a result. At configuration $\sigma(s)$, with $k+1 \leq s \leq i_0 - 1$, from 111 to 162 graph rules are used, depending of the length of the history being built at the beginning of the run; once the length of history has reach its maximum, i.e., $|\mathcal{M}| + 1$, the most used graph rules are the ones designed to clear outdated history. At configuration $\sigma(i_0)$, the reconfiguration **removegps** occurs, then as long as the event **addgps normal** has not occurred yet, only the evaluation of the part “... **before addgps normal**” of the formula is needed to obtain a result; 237 graph rules are used, most of them doing a cleaning of the elements of the subgraph representing the monitor of the GPS component being removed. At configuration $\sigma(s)$, with $k+1 \leq s \leq i_0 - 1$, 149 graph rules are used, mainly to clear outdated history. At configuration $\sigma(i_1)$,

Table 1: Graph rules used at configuration $\sigma(s)$

Index s of configuration	Number of graph rules	Reconfiguration	Part of formula to be evaluated
$s = k$	85		after removegps normal ...
$k + 1 \leq s \leq i_0 - 1$	111 – 162		after removegps normal ...
$s = i_0$	237	removegps	... before addgps normal
$i_0 + 1 \leq s \leq i_1 - 1$	149		... before addgps normal
$s = i_1$	253	addgps	... eventually ($\text{power} \geq 33$)...

the reconfiguration `addgps` occurs, then only the evaluation, at the previous configuration, of the part “...**eventually** ($power \geq 33$)...” of the formula is needed to obtain a result; 253 graph rules are used, mainly to clear outdated history and to update the scope of the property.

6 Conclusion

This paper has addressed the decentralised evaluation problem for linear temporal patterns on reconfiguration paths of component-based systems. To this end, we have proposed a specific progressive semantics of temporal patterns, and an algorithm for their decentralised evaluation using monitors associated with components. We have shown that when reached, the decentralised evaluation results coincide with the results obtained by the centralised evaluation of temporal patterns at runtime. We have described the implementation with *GROOVE* and its application to a location composite component.

In this paper, for the sake of readability, monitors only deal with a single FTPL property. To evaluate several FTPL formulae, we can either use a single monitor (per component) dealing with all the formulae, as herein described, or a monitor per formula of interest. Depending on the context, each method can have its own advantages and drawbacks.

In the case of the removal of a component, the corresponding monitor terminates and is removed. Thanks to the adaptation policies’ controller, this should not influence any ongoing temporal pattern evaluation. When a component is added, its monitor starts with a blank history. Furthermore, when a monitored primitive component is replaced with a composite component whose sub-components contain (among other) the same parameters as the original component, the monitor shall keep working seamlessly. Since no additional monitor is added, this mechanism allows us to mitigate the communication overhead that could be incurred by the increase of the number of components.

As a future work, we intend to extend the analysis of the TPDEP problem to the case when several reconfiguration operations occur. It would be possible when reconfigurations lead to configurations whose atomic events do not interfere with the evaluation of the temporal property of interest (the TPDEP input). In this case the adaptation controller can authorize reconfigurations of independent parts of the component-based system. On the implementation side, we plan to exploit the decentralised evaluation method for the implementation handling the adaptation policies. The overall goal is to exploit its results to apply adaptation policies to the component-based system under scrutiny at runtime. So far we have considered the components having monitors to be all on the same architectural level, i.e., they all are siblings. As a future work, we plan to *delegate* part of the monitoring of composite components to their subcomponents.

References

1. de Lemos, R., et. al: Software engineering for self-adaptive systems: A second research roadmap. In de Lemos, R., Giese, H., Müller, H.A., Shaw, M., eds.: *Software Engineering for Self-Adaptive Systems*. Volume 7475 of LNCS., Springer (2010) 1–32
2. Dwyer, M.B., Avrunin, G.S., Corbett, J.C.: Patterns in property specifications for finite-state verification. In: *Software Engineering*. (1999) 411–420
3. Kouchnarenko, O., Weber, J.F.: Adapting component-based systems at runtime via policies with temporal patterns. In Fiadeiro, J.L., Liu, Z., Xue, J., eds.: *FACS*. Volume 8348 of LNCS. Springer (2014) 234–253
4. Bauer, A., Falcone, Y.: Decentralised LTL monitoring. In: *FM 2012: Formal Methods*. Volume 7436 of LNCS., Springer (2012) 85–100
5. Ghamarian, A.H., de Mol, M., Rensink, A., Zambon, E., Zimakova, M.: Modelling and analysis using GROOVE. *Int. J. on Software Tools for Technology Transfer* **14** (2012) 15–40
6. Kupferman, O., Vardi, M.Y., Wolper, P.: Module checking. *Information and Computation* **164** (2001) 322–344
7. Bozzelli, L., Murano, A., Peron, A.: Pushdown module checking. In Sutcliffe, G., Voronkov, A., eds.: *LPAR*. Volume 3835 of LNCS., Springer (2005) 504–518
8. Aminof, B., Legay, A., Murano, A., Serre, O., Vardi, M.Y.: Pushdown module checking with imperfect information. *Inf. Comput.* **223** (2013) 1–17
9. Bauer, A., Leucker, M., Schallhart, C.: Comparing LTL semantics for runtime verification. *Int. J. of Logic and Computation* **20** (2010) 651–674
10. Kim, M., Lee, I., Shin, J., Sokolsky, O., et al.: Monitoring, checking, and steering of real-time systems. *ENTCS* **70** (2002) 95–111
11. Bruneton, E., Coupaye, T., Leclercq, M., Quéma, V., Stefani, J.B.: The fractal component model and its support in java. *Softw., Pract. Exper.* **36** (2006) 1257–1284
12. Schneider, S., Treharne, H.: CSP theorems for communicating B machines. *Formal Asp. Comput.* **17** (2005) 390–422
13. Seinturier, L., Merle, P., Rouvoy, R., Romero, D., Schiavoni, V., Stefani, J.B.: A component-based middleware platform for reconfigurable service-oriented architectures. *Software: Practice and Experience* **42** (2012) 559–583
14. Garavel, H., Mateescu, R., Serwe, W.: Large-scale distributed verification using cadp: Beyond clusters to grids. *Electr. Notes Theor. Comput. Sci.* **296** (2013) 145–161
15. Dormoy, J., Kouchnarenko, O., Lanoix, A.: Runtime verification of temporal patterns for dynamic reconfigurations of components. In Arbab, F., Ölveczky, P., eds.: *FACS*. Volume 7253 of LNCS. Springer Berlin Heidelberg (2012) 115–132
16. Dormoy, J., Kouchnarenko, O., Lanoix, A.: Using temporal logic for dynamic reconfigurations of components. In Barbosa, L., Lumpe, M., eds.: *FACS*. Volume 6921 of LNCS. Springer Berlin Heidelberg (2012) 200–217
17. Hamilton, A.G.: *Logic for mathematicians*. Cambridge University Press, Cambridge (1978)
18. Dormoy, J., Kouchnarenko, O., Lanoix, A.: When structural refinement of components keeps temporal properties over reconfigurations. In: *FM 2012: Formal Methods*. Volume 7436 of LNCS. Springer (2012) 171–186

A FTPL Syntax and Semantics [3]

In this section, we summarize the FTPL semantics as presented in [3].

A.1 Syntax

$\langle FTPL \rangle$	$::= \langle tpp \rangle \mid \langle events \rangle \mid cp$
$\langle tpp \rangle$	$::= \textbf{after} \langle events \rangle \langle tpp \rangle \mid \textbf{before} \langle events \rangle \langle trp \rangle \mid \langle trp \rangle \textbf{until} \langle events \rangle \mid \langle trp \rangle$
$\langle trp \rangle$	$::= \textbf{always} cp \mid \textbf{eventually} cp \mid \langle trp \rangle \wedge \langle trp \rangle \mid \langle trp \rangle \vee \langle trp \rangle$
$\langle events \rangle$	$::= \langle event \rangle, \langle events \rangle \mid \langle event \rangle$
$\langle event \rangle$	$::= ope \textbf{normal} \mid ope \textbf{exceptional} \mid ope \textbf{terminates} \mid ext$

A.2 FTPL Basic Semantics

Definition 10 (Configuration Properties Semantics). We say that a configuration $\sigma(i)$ satisfies $cp \in CP$, written $\llbracket \sigma(i) \models cp \rrbracket = \top$, when $l(\sigma(i)) \Rightarrow cp$. In this case, cp is valid on $\sigma(i)$. Otherwise, $\sigma(i)$ does not satisfy cp , written $\llbracket \sigma(i) \models cp \rrbracket = \perp$.

Definition 11 (FTPL Events Semantics). Let ope be a reconfiguration operation, ext an external event, e an event, and $events$ a list of events.

The interpretation of the events at the i -th state of the path σ is defined by:

$$\begin{aligned}
\llbracket \sigma(i) \models ope \textbf{normal} \rrbracket &= \begin{cases} \top & \text{if } i > 0 \wedge \sigma(i-1) \neq \sigma(i) \wedge \sigma(i-1) \xrightarrow{ope} \sigma(i) \\ \perp & \text{otherwise.} \end{cases} \\
\llbracket \sigma(i) \models ope \textbf{exceptional} \rrbracket &= \begin{cases} \top & \text{if } i > 0 \wedge \sigma(i-1) = \sigma(i) \wedge \sigma(i-1) \xrightarrow{ope} \sigma(i) \\ \perp & \text{otherwise.} \end{cases} \\
\llbracket \sigma(i) \models ope \textbf{terminates} \rrbracket &= \llbracket \sigma(i) \models ope \textbf{normal} \rrbracket \sqcup \llbracket \sigma(i) \models ope \textbf{exceptional} \rrbracket \\
\llbracket \sigma(i) \models ext \rrbracket &= eval_{\sigma}(cp_{ext}, i)^{12} \\
\llbracket \sigma(i) \models e, events \rrbracket &= \llbracket \sigma(i) \models e \rrbracket \sqcup \llbracket \sigma(i) \models events \rrbracket
\end{aligned}$$

A.3 FTPL Runtime Progressive Semantics

In order to evaluate, in a *progressive* fashion, FTPL expressions at runtime, without consulting a complete history of FTPL properties' evaluation, we introduce the following notations. Let $\phi_{\sigma} = \llbracket \sigma \models \phi \rrbracket$ be the evaluation of an FTPL formula where ϕ is a list of events, a trace property, or a temporal property. We denote $\phi_{\sigma}(i)$ the evaluation of ϕ on σ , at the i -th state of the path.

Furthermore, following [2], if the scope of an FTPL property ϕ is restricted to the suffix path σ_k , $k \geq 0$, we write $\phi_{\sigma_k} = \llbracket \sigma_k \models \phi \rrbracket$ for such a restriction, and $\phi_{\sigma_k}(i)$ for the evaluation in \mathbb{B}_4 of this restriction at the i -th state of σ , where $i \geq k$. Then, the evaluation of ϕ on the path σ ($\phi_{\sigma} = \llbracket \sigma \models \phi \rrbracket$), is similar to the evaluation of ϕ on the suffix path σ_0 starting at the first configuration, which is $\phi_{\sigma_0} = \llbracket \sigma_0 \models \phi \rrbracket$. For the sake of simplicity, we also write $cp_{\sigma_k}(i) = \llbracket \sigma_k(i) \models cp \rrbracket$.

¹² ext , $eval_{\sigma}$, and cp_{ext} are defined as in [3].

Definition 12 (FTPL Runtime Progressive Trace Properties Semantics). Let cp be a configuration property, ϕ (resp. φ) a trace property of the form $\phi = \mathbf{always} \ cp$ (resp. $\varphi = \mathbf{eventually} \ cp$). We define $\phi_{\sigma_k}(i)$ (resp. $\varphi_{\sigma_k}(i)$), the evaluation in \mathbb{B}_4 of $\llbracket \sigma_k \models \phi \rrbracket$ (resp. $\llbracket \sigma_k \models \varphi \rrbracket$) at the i -th state of σ when the scope is restricted to σ_k , by:

- for $i = k$, $\phi_{\sigma_k}(k) = \top^p \sqcap cp_{\sigma}(k)$; $\varphi_{\sigma_k}(k) = \perp^p \sqcup cp_{\sigma}(k)$
- for $i > k$, $\phi_{\sigma_k}(i) = \phi_{\sigma_k}(i-1) \sqcap cp_{\sigma}(i)$; $\varphi_{\sigma_k}(i) = \varphi_{\sigma_k}(i-1) \sqcup cp_{\sigma}(i)$

Furthermore, let ψ_1 and ψ_2 be two trace properties, then:

$$\llbracket \sigma_k \models \psi_1 \wedge \psi_2 \rrbracket = \llbracket \sigma_k \models \psi_1 \rrbracket \sqcap \llbracket \sigma_k \models \psi_2 \rrbracket ; \llbracket \sigma_k \models \psi_1 \vee \psi_2 \rrbracket = \llbracket \sigma_k \models \psi_1 \rrbracket \sqcup \llbracket \sigma_k \models \psi_2 \rrbracket$$

On the scope starting at the k -th state of σ , if at state k one has $cp_{\sigma}(k) = \top$ (resp. $cp_{\sigma}(k) = \perp$), the trace property **always** cp (resp. **eventually** cp) is evaluated to \top^p (resp. \perp^p); otherwise, it is evaluated to \perp (resp. \top). Then, for $i > k$, at the i -th state of σ , **always** cp (resp. **eventually** cp) is evaluated to the minimum (resp. maximum), interpreted wrt. \sqsubseteq , of a) its evaluation at the previous state and b) $cp_{\sigma}(i)$.

Definition 13 (FTPL Runtime Progressive Lists of Events Semantics). Let e be a list of events. We define $e_{\sigma_k}(i)$, the evaluation in \mathbb{B}_4 of $\llbracket \sigma_k \models e \rrbracket$ at the i -th state of σ when the scope is restricted to σ_k , by:

- for $i = k$, $e_{\sigma_k}(k) = \llbracket \sigma_k(k) \models e \rrbracket$
- for $i > k$, $e_{\sigma_k}(i) = \llbracket \sigma_k(i) \models e \rrbracket \sqcup (\top^p \sqcap e_{\sigma_k}(i-1))$

Intuitively, the expression $\llbracket \sigma(i) \models e \rrbracket \sqcup (\top^p \sqcap e_{\sigma_k}(i-1))$ evaluates to \top if there is an occurrence of e at configuration i , and to \perp (resp. \top^p) if there is no occurrence of e at configuration i and no (resp. at least one) occurrence of e happening before configuration i on the scope starting at configuration k .

Definition 14 (FTPL Runtime Progressive Temporal Properties Semantics). Let tpp be a temporal property, trp a trace property, e a list of events, ϕ (resp. φ, ψ) a temporal property of the form $\phi = \mathbf{after} \ e \ tpp$ (resp. $\varphi = \mathbf{before} \ e \ trp, \psi = trp \ \mathbf{until} \ e$). We define $\phi_{\sigma_k}(i)$ (resp. $\varphi_{\sigma_k}(i), \psi_{\sigma_k}(i)$), the evaluation in \mathbb{B}_4 of $\llbracket \sigma_k \models \phi \rrbracket$ (resp. $\llbracket \sigma_k \models \varphi \rrbracket, \llbracket \sigma_k \models \psi \rrbracket$) at the i -th state of σ when the scope is restricted to σ_k , by: for $i \geq k$,

$$\phi_{\sigma_k}(i) = \left(\bigcap_{j \in \mathcal{I}_{\sigma_k}^i(e)} tpp_{\sigma_j}(i) \right) \sqcap \top^p \quad \text{where } \mathcal{I}_{\sigma_k}^i(e) = \{j | k \leq j \leq i \wedge \llbracket \sigma(j) \models e \rrbracket = \top\}$$

represents the set of indexes for an occurrence of e .

$$\varphi_{\sigma_k}(i) = \begin{cases} \top^p & \text{if } e_{\sigma_k}(i) = \perp \vee i = k \\ \perp & \text{if } e_{\sigma_k}(i) = \top \wedge trp_{\sigma_k}(i-1) \in \{\perp, \perp^p\} \\ \varphi_{\sigma_k}(i-1) & \text{otherwise} \end{cases}$$

$$\psi_{\sigma_k}(i) = \begin{cases} \top^p & \text{if } trp_{\sigma_k}(i) \neq \perp \wedge e_{\sigma_k}(i) = \top \wedge e_{\sigma_k}(i-1) = \perp \wedge trp_{\sigma_k}(i-1) \in \{\top^p, \top\} \\ \perp^p & \text{if } trp_{\sigma_k}(i) \neq \perp \wedge (e_{\sigma_k}(i) = \perp \vee i = k) \\ \perp & \text{if } trp_{\sigma_k}(i) = \perp \vee (e_{\sigma_k}(i) = \top \wedge trp_{\sigma_k}(i-1) \in \{\perp, \perp^p\}) \\ \psi_{\sigma_k}(i-1) & \text{otherwise} \end{cases}$$

By definition, the evaluation of $\phi = \mathbf{after} \ e \ tpp$ is either $a)$ \top^p as long as e does not occur or if tpp is evaluated to \top^p or \top on each suffix of the path

starting at an occurrence of e , or b) \perp if on any of these suffixes tpp is evaluated to \perp , or c) \perp^p , otherwise.

For $\varphi = \mathbf{before} \ e \ trp$, its evaluation is either a) \top^p if e has not occurred yet, or b) \perp if for each occurrence of e , trp is evaluated to \perp or \perp^p on the segment starting at the beginning of the considered scope and ending at the previous $i - 1$ -th configuration on the σ path. Otherwise, ϕ at the i -th configuration is evaluated to its value at the previous $i - 1$ -th configuration.

Intuitively, the $\psi = trp \ \mathbf{until} \ e$ property can be seen as being evaluated similarly to $\mathbf{before} \ e \ trp$, but with the two following exceptions: a) when trp is evaluated to \perp , ψ is evaluated to \perp ; otherwise, b) on the beginning part of the scope and as long as e has not occurred, ψ is evaluated to \perp^p .

Finally, we say that a reconfiguration model $S = \langle \mathcal{C}, \mathcal{C}^0, \mathcal{R}_{run}, \rightarrow, l \rangle$ satisfies a property $\phi \in Prop_{FTPL}$, written $S \models \phi$, if $\forall \sigma. (\sigma \in \Sigma(S) \wedge \sigma(0) \in \mathcal{C}^0 \Rightarrow \phi_\sigma = \top)$.

B Progression of FTPL Properties

Let us introduce the formulae in Fig. 8 to ease the notations of the progression function for FTPL properties. In Equation 6, $\mathcal{I}_{\sigma_k}(e)$ is the set of indexes of the segment σ_k where the FTPL event e is evaluated to \top . Equation 7 represents a rewriting formula that associates the progression of an FTPL property with a couple composed of itself and of the progression of this property at the previous event. Equations 8 to 10 are established using the FTPL temporal properties' semantics.

With the formulae of Fig. 8, the progression function P of Def. 4 is extended to the FTPL properties' evaluation formulae on a (suffix) path as follows.

$$F_{\mathcal{I}}(\mathcal{I}_{\sigma_k}(e), \theta(i)) = \begin{cases} F_{\mathcal{I}}(\mathcal{I}_{\sigma_k}(e), \theta(i-1)) \cup \{j | j = i \wedge P(e, \theta(j)) = \top\} & \text{for } i \geq k \\ \emptyset & \text{otherwise} \end{cases} \quad (6)$$

$$F_{\mathcal{P}}(\xi, \theta(i)) = (P(\xi, \theta(i)), P(\xi, \theta(i-1))) \quad (7)$$

$$F_{\mathcal{A}}(\mathcal{J}, P(e, \theta(i)), P(tpp_{\sigma_i}, \theta(i))) = \top^p \sqcap \prod_{j \in \mathcal{J}} P(tpp_{\sigma_j}, \theta(i)) \quad (8)$$

$$F_{\mathcal{B}}(\varepsilon, trp, tpp) = \begin{cases} \top^p & \text{if } \varepsilon = \perp \\ \perp & \text{if } \varepsilon = \top \wedge trp \in \{\perp, \perp^p\} \\ tpp & \text{otherwise} \end{cases} \quad (9)$$

$$F_{\mathcal{U}}((\varepsilon, \varepsilon'), (trp, trp'), tpp) = \begin{cases} \top^p & \text{if } trp \neq \perp \wedge \varepsilon = \top \wedge \varepsilon' = \perp \wedge trp' \in \{\top^p, \top\} \\ \perp^p & \text{if } trp \neq \perp \wedge \varepsilon = \perp \\ \perp & \text{if } trp = \perp \vee (\varepsilon = \top \wedge trp' \in \{\perp, \perp^p\}) \\ tpp & \text{otherwise} \end{cases} \quad (10)$$

Fig. 8: Rewriting formulae for progression of FTPL properties

Definition 15 (Progression of FTPL properties' evaluation formulae on a (suffix) path). Let cp be a configuration property, ψ_1 and ψ_2 two trace properties, ϕ (resp. φ) a trace property of the form $\phi = \mathbf{always}$ cp (resp. $\varphi = \mathbf{eventually}$ cp), e a list of FTPL events, tpp a temporal property, trp a trace property, α (resp. β , γ) a temporal property of the form $\alpha = \mathbf{after}$ e tpp (resp. $\beta = \mathbf{before}$ e trp , $\gamma = trp$ **until** e). The progression function P for FTPL properties on a (suffix) path is defined by:

$$P(\alpha_{\sigma_k}, \theta(i)) = F_A(F_I(\mathcal{I}_{\sigma_k}(e), \theta(i)), P(e, \theta(i)), P(tpp_{\sigma_i}, \theta(i))) \text{ for } i \geq k \quad (11)$$

$$P(\beta_{\sigma_k}, \theta(i)) = \begin{cases} \top^P & \text{for } i = k \\ F_B(P(e_{\sigma_k}, \theta(i)), P(\overline{\mathbf{X}}trp_{\sigma_k}, \theta(i)), P(\overline{\mathbf{X}}\beta_{\sigma_k}, \theta(i))) & \text{for } i > k \end{cases} \quad (12)$$

$$P(\gamma_{\sigma_k}, \theta(i)) = \begin{cases} P(trp_{\sigma_k}, \theta(i)) \sqcap \perp^P & \text{for } i = k \\ F_U(F_P(e_{\sigma_k}, \theta(i)), F_P(trp_{\sigma_k}, \theta(i)), P(\overline{\mathbf{X}}\gamma_{\sigma_k}, \theta(i))) & \text{for } i > k \end{cases} \quad (13)$$

$$P(e_{\sigma_k}, \theta(i)) = \begin{cases} P(e, \theta(i)) & \text{for } i = k \\ P(e, \theta(i)) \sqcup (\top^P \sqcap P(\overline{\mathbf{X}}e_{\sigma_k}, \theta(i))) & \text{for } i > k \end{cases} \quad (14)$$

$$P(\phi_{\sigma_k}, \theta(i)) = \begin{cases} P(cp, \theta(i)) \sqcap \top^P & \text{for } i = k \\ P(cp, \theta(i)) \sqcap P(\overline{\mathbf{X}}\phi_{\sigma_k}, \theta(i)) & \text{for } i > k \end{cases} \quad (15)$$

$$P(\varphi_{\sigma_k}, \theta(i)) = \begin{cases} P(cp, \theta(i)) \sqcup \perp^P & \text{for } i = k \\ P(cp, \theta(i)) \sqcup P(\overline{\mathbf{X}}\varphi_{\sigma_k}, \theta(i)) & \text{for } i > k \end{cases} \quad (16)$$

$$P(\psi_1 \wedge \psi_2, \theta(i)) = P(\psi_1, \theta(i)) \sqcap P(\psi_2, \theta(i)) \quad (17)$$

$$P(\psi_1 \vee \psi_2, \theta(i)) = P(\psi_1, \theta(i)) \sqcup P(\psi_2, \theta(i)) \quad (18)$$