



HAL
open science

Practical Analysis Framework for Component Systems with Dynamic Reconfigurations

Olga Kouchnarenko, Jean-François Weber

► **To cite this version:**

Olga Kouchnarenko, Jean-François Weber. Practical Analysis Framework for Component Systems with Dynamic Reconfigurations. 17th International Conference on Formal Engineering Methods, Nov 2015, Paris, France. hal-01135720v3

HAL Id: hal-01135720

<https://hal.science/hal-01135720v3>

Submitted on 15 Jul 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Practical Analysis Framework for Component Systems with Dynamic Reconfigurations^{*}

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

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

² Inria/Nancy-Grand Est, Villers-lès-Nancy, France
{okouchnarenko,jfweber}@femto-st.fr

Abstract. Dynamic reconfigurations that modify the architecture of component-based systems without incurring any system downtime need to preserve the architectural consistency. In this context, we propose a reconfiguration model based on Hoare logic using sequences and (unlike most of the related work on reconfigurations) the alternative and the repetitive constructs. Using primitive reconfiguration operations as building blocks, this model takes advantage of the predicate-based semantics of programming language constructs and weakest preconditions to treat dynamic reconfigurations in a manner that preserves configuration consistency. Then, after enriching the model with interpreted configurations and reconfigurations in a consistency compatible manner, a conformance relation is exploited to validate component systems' implementations within the environment supporting the Fractal and FraSCAti frameworks. A practical contribution consists of promising experimental results obtained thanks to our implementations, notably on a cloud-based multi-tier hosting environment model managed as a component system.

1 Introduction

Dynamic reconfigurations that modify the architecture of self-adaptive [1] component-based systems without incurring any system downtime must happen not only in suitable circumstances, but also need to preserve the architectural consistency. Whereas the former can be ensured by adaptation policies [1,2], the latter is directly related to the definition of reconfigurations and to the reconfiguration ordering/protocol [3,4].

In [3], it is assumed that the reconfigurations always make the component assembly evolve from one consistent architecture to another consistent architecture, only through a path of architecturally consistent architectures. However, primitive reconfigurations like *unbind*, *stop*, etc. may disrupt such a path. With relation to consistency constraints defined in [5] over component-based architectures, their preservation of the system under scrutiny was uneasy to prove, mostly because of the lack of precise semantics for primitive reconfiguration operations. Therefore, when considering more complicated reconfigurations composed of sequences,

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

repetitions, or choices over primitive reconfiguration operations, to address the above-mentioned issue, we propose to express reconfigurations' preconditions and postconditions using the concept of *weakest precondition* [6]. This precise and concise formalism allows us to express primitive and non primitive *guarded reconfigurations*; this is the first contribution simplifying both reconfiguration protocols and adaptation policies.

Then, after enriching the model with interpreted configurations and reconfigurations in a consistency-compatible way, a conformance relation is exploited to validate implementations of a component architectural model developed within our architecture manager supporting the Fractal [7] and FraSCAti [8] frameworks.

This second practical contribution allows us, not only, to simulate a desired run of a system being reconfigured, but, also, to generate all (or a subset of the) possible reconfiguration combinations useful, for example, for a (bounded) reachability analysis. The paper reports on promising experimental results obtained thanks to our implementations, notably on a cloud-based multi-tier application hosting environment model managed as a component software architecture.

The paper is organised as follows: Section 2 presents, as a case study, a cloud-based multi-tier application hosting environment managed as a component-based system. Background information on our component-based reconfiguration model, as well as elements of operational semantics are given in Sect. 3. In Sect. 4 a richer interpreted reconfiguration model is shown to be weakly simulated by the more abstract model; nevertheless, this simulation respects non-divergency. Using several case studies, Sect. 5 describes conformant implementations of the interpreted model within different environments. Section 6 presents related work and our conclusion.

2 Case Study

Internet service providers and telecommunications operators tend more and more to define themselves as cloud providers. In this context, automation of software and (virtual) hardware installation and configuration is paramount. It is not enough for an application to be cloud-ready; it has to be scalable and scalability mechanisms need to be integrated in the core of the cloud management system.

We consider a typical three-tier web application using a front-end web server, a middle-ware application server, and a back-end data providing service such as a database or a data store. Figure 1 shows a single virtual machine (or *VM*) hosting together the three services of such an application.

The VM is represented as a composite component *virtualMachine* containing sub-components representing each service (*httpServer*, *appServer*, and *dataServer*) of the application. Each of the service sub-component has two provided interfaces: one to provide its service and another one used to monitor the service.




Fig. 1: Managed Virtual Machine with Three-tier Application Components

Furthermore, the VM of Fig. 1 also contains four *observers*, that are sub-components used to monitor services. The sub-component *osObs* is used to monitor the Operating System of the VM. It is also bound to the sub-components *httpObs*, *appObs*, and *dataObs* used respectively to monitor the services of the *httpServer*, *appServer*, and *dataServer* sub-components. Finally, the VM composite component itself has two provided interfaces: one used to provide services and a second one used for monitoring.

Of course, a VM does not have to be monitored, nor have to host the three types of services. Figure 2 illustrates a *cloud environment*, *clouEnv*, containing a VM used for development purpose (*vmDev*) that contains the three tiers of the application without being monitored; such a VM is called *unmanaged*. The three other VM are all monitored, i.e., *managed*, and each contains a tier of the application. The reader can note that each of the managed VM contains only the observers responsible for monitoring the operating system and the type of service provided. The cloud environment has three provided interfaces: two to provide its service, whether it is or not in a development version, and another one, used for monitoring, connected to a sub-component *monitorObs* bound to all the monitoring interfaces of the managed VM.

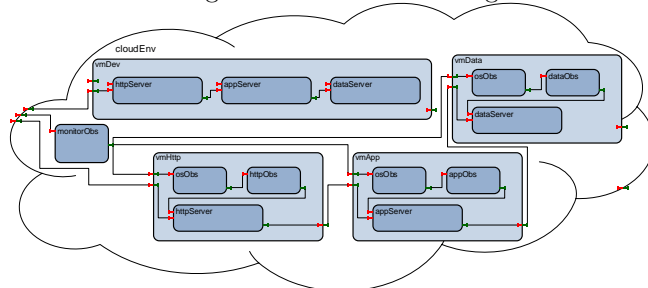


Fig. 2: Cloud Environment Example

Depending on the services to provide and the monitoring state (managed vs unmanaged) the necessary components should be added. During the life cycle of the VM some configuration changes can happen; we consider them as reconfigurations of a component-based system. In this context, the challenge consists in performing adequate dynamic reconfigurations with minimum communication overhead, while avoiding reconfigurations that would lead to unwanted behaviours.

A cloud provider must be able to provide on-demand (sets of) VMs configured with the right service components and the appropriate monitoring. In this context, we study the provisioning of a single VM as illustrated in Fig. 1.

3 Component-based Architecture

3.1 Configurations and Reconfigurations

Component models can be very heterogeneous. Most of them consider software components that can be seen as black boxes (or grey boxes if some of their inner features are visible) having fully-described interfaces. Behaviours and interactions are specified using components' definitions and their interfaces. In this section, we revisit the architectural reconfiguration model introduced in [5,9]. In general, the system configuration is the specific definition of the elements that define or

prescribe what a system is composed of, while a reconfiguration can be seen as a transition from a configuration to another.

Following [9], a configuration is defined to be a set of architectural elements (components, required or provided interfaces, and parameters) together with relations to structure and to link them.

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;*
 - *$Interfaces = RequiredInts \uplus ProvidedInts$ 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;*
- *$Rel = \left\{ \begin{array}{l} Container \uplus ContainerType \uplus Contingency \\ \uplus Parent \uplus Binding \uplus Delegate \uplus State \uplus Value \end{array} \right.$* 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 which 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 to each (required or provided) interface and to each parameter;*
 - *$Contingency : RequiredInts \rightarrow \{\text{mandatory}, \text{optional}\}$ is a total function indicating whether each required interface is mandatory or optional;*
 - *$Parent \subseteq Components \times Components$ is a relation linking a sub-component to the corresponding composite component³;*
 - *$Binding : ProvidedInts \rightarrow RequiredInts$ is a partial function which binds together a provided interface and a required one;*
 - *$Delegate : Interfaces \rightarrow Interfaces$ is a partial function to express delegation links;*
 - *$State : Components \rightarrow \{\text{started}, \text{stopped}\}$ is a total function giving the status of instantiated components;*
 - *$Value : Parameters \rightarrow \{t | t \in PType\}$ is a total function which gives the current value (of type $t \in PType$) of each parameter.*

We also introduce a set CP of configuration propositions which are constraints on the architectural elements and the relations between them. These propositions are specified using first-order logic formulae [10]. The interpretation of functions, relations, and predicates over $Elem$ is done according to basic definitions in [10] and Def. 1. The interested reader is referred to [5].

³ For any $(p, q) \in Parent$, we say that q has a sub-component p , i.e. p is a child of q . Shared components (sub-components of multiple enclosing composite components) can have more than one parent.

Let $\mathcal{C} = \{c, c_1, c_2, \dots\}$ be a set of configurations. An *interpretation* function $l : \mathcal{C} \rightarrow CP$ gives the largest conjunction of $cp \in CP$ evaluated to true on $c \in \mathcal{C}$. We say that a configuration $c = \langle Elem, Rel \rangle$ satisfies $cp \in CP$, when $l(c) \Rightarrow cp$; in this case, cp is valid on c , otherwise, c does not satisfy cp .

Among the configuration propositions, the architectural *consistency constraints* CC in Table 1 express requirements on component assembly common to all the component architectures [5]. Intuitively,

- a component *supplies*, at least, one provided interface (**CC.1**);
- the composite components have no parameter (**CC.2**);
- a sub-component must not include its own parent component (**CC.3**);
- two bound interfaces must have the same interface type (**CC.4**) and their containers are sub-components of the same composite (**CC.5**);
- when binding two interfaces, there is a need to ensure that they have not been involved in a delegation yet (**CC.6**); similarly, when establishing a delegation link between two interfaces, the specifier must ensure that they have not yet been involved in a binding (**CC.7**);
- a provided (resp. required) interface of a sub-component is delegated to at most one provided (resp. required) interface of its parent component (**CC.8**), (**CC.9**) and (**CC.11**); the interfaces involved in the delegation must have the same interface type (**CC.10**); and
- a component is *started* only if its mandatory required interfaces are bound or delegated (**CC.12**).

Table 1: Consistency Constraints

$$\begin{array}{l}
 \forall c. (c \in \text{Components} \Rightarrow (\exists ip. (ip \in \text{ProvidedInts} \wedge \text{Container}(ip) = c))) \quad (\text{CC.1}) \\
 \forall c, c' \in \text{Components}. ((c, c') \in \text{Parent} \Rightarrow \forall p. (p \in \text{Parameters} \Rightarrow \text{Container}(p) \neq c')) \quad (\text{CC.2}) \\
 \forall c, c' \in \text{Components}. ((c, c') \in \text{Parent}^+ \Rightarrow c \neq c') \quad (\text{CC.3}) \\
 \forall ip \in \text{ProvidedInts}, \forall ir \in \text{RequiredInts}. \left(\text{Binding}(ip) = ir \Rightarrow \begin{array}{l} \text{ContainerType}(ip) = \text{ContainerType}(ir) \\ \wedge \text{Container}(ip) \neq \text{Container}(ir) \end{array} \right) \quad (\text{CC.4}) \\
 \forall ip \in \text{ProvidedInts}, \forall ir \in \text{RequiredInts}. \left(\text{Binding}(ip) = ir \Rightarrow \exists c \in \text{Components}. \left(\begin{array}{l} (\text{Container}(ip), c) \in \text{Parent} \\ \wedge (\text{Container}(ir), c) \in \text{Parent} \end{array} \right) \right) \quad (\text{CC.5}) \\
 \forall ip \in \text{ProvidedInts}, \forall ir \in \text{RequiredInts}. \left(\text{Binding}(ip) = ir \Rightarrow \begin{array}{l} ip \notin \text{dom}(\text{Delegate}) \\ \wedge ir \notin \text{dom}(\text{Delegate}) \end{array} \right) \quad (\text{CC.6}) \\
 \forall i, i' \in \text{Interfaces}. (\text{Delegate}(i) = i' \Rightarrow i \notin \text{dom}(\text{Binding}) \wedge i \notin \text{codom}(\text{Binding})) \quad (\text{CC.7}) \\
 \forall i, i' \in \text{Interfaces}. (\text{Delegate}(i) = i' \wedge i \in \text{ProvidedInts} \Rightarrow i' \in \text{ProvidedInts}) \quad (\text{CC.8}) \\
 \forall i, i' \in \text{Interfaces}. (\text{Delegate}(i) = i' \wedge i \in \text{RequiredInts} \Rightarrow i' \in \text{RequiredInts}) \quad (\text{CC.9}) \\
 \forall i, i' \in \text{Interfaces}. \left(\text{Delegate}(i) = i' \Rightarrow \begin{array}{l} \text{ContainerType}(i) = \text{ContainerType}(i') \\ \wedge (\text{Container}(i), \text{Container}(i')) \in \text{Parent} \end{array} \right) \quad (\text{CC.10}) \\
 \forall i, i', i'' \in \text{Interfaces}. \left(\begin{array}{l} (\text{Delegate}(i) = i' \wedge \text{Delegate}(i) = i'' \Rightarrow i' = i'') \\ \wedge (\text{Delegate}(i) = i'' \wedge \text{Delegate}(i') = i'' \Rightarrow i = i') \end{array} \right) \quad (\text{CC.11}) \\
 \forall ir \in \text{RequiredInts}. \left(\begin{array}{l} \text{State}(\text{Container}(ir)) = \text{started} \\ \wedge \text{Contingency}(ir) = \text{mandatory} \end{array} \Rightarrow \exists i \in \text{Interfaces}. \left(\begin{array}{l} \text{Binding}(i) = ir \\ \vee \text{Delegate}(ir) = i \end{array} \right) \right) \quad (\text{CC.12})
 \end{array}$$

Definition 2 (Consistent configuration). Let $c = \langle Elem, Rel \rangle$ be a configuration and CC the consistency constraints. The configuration c is consistent, written $\text{consistent}(c)$, if $l(c) \Rightarrow CC$. We write $\text{consistent}(\mathcal{C})$ when $\forall c \in \mathcal{C}. \text{consistent}(c)$.

3.2 Reconfiguration Model and Consistency Propagation

Reconfigurations make the component-based architecture evolve dynamically. They are composed of primitive operations such as instantiation/destruction (*new/destroy*) of components; addition/removal (*add/remove*) of components; binding/unbinding (*bind/unbind*) of component interfaces; starting/stopping (*start/stop*) components; setting components' parameters values (*update*). These primitive operations obey pre/post predicates. For example, before adding a sub-component $comp_1$ to a composite $comp_2$, one must verify, as in Table 2, that *a*) $comp_1$ and $comp_2$ exist (2) and are different (3), *b*) $comp_2$ is not a descendant of $comp_1$ (4), and *c*) $comp_2$ has no parameter (5). When these preconditions are met, the postcondition consists in adding $(comp_1, comp_2)$ to the *Parent* relation, as expressed by $R_{add} = Parent \cup \{(comp_1, comp_2)\}$ (1).

Table 2: Preconditions of the *add* primitive reconfiguration operation

$$\begin{array}{ll}
 comp_1, comp_2 \in Components & (2) \qquad (comp_2, comp_1) \notin Parent^+ & (4) \\
 comp_1 \neq comp_2 & (3) \quad \forall p \in Parameters.Container(p) \neq comp_2 & (5)
 \end{array}$$

Inspired by the predicate-based semantics of programming language constructs [11], we consider a reconfiguration operation *ope*, and two configurations c and c' such that the transition between c and c' is performed using *ope* (denoted by $c \xrightarrow{ope} c'$). Then, given R , some conditions on the configuration of the system under scrutiny, the notation $wp(ope, R)$ denotes, as in [6], the *weakest precondition* for the configuration c such that activation of *ope* can occur and, if so, is guaranteed to lead to c' satisfying the postcondition R . More formally, in our case, if $l(c) \Rightarrow wp(ope, R)$ and $c \xrightarrow{ope} c'$ then $l(c') \Rightarrow R$. Therefore, considering the *add* primitive reconfiguration operation whose preconditions are displayed in Table 2, the weakest precondition $wp(add, R_{add})$ is the conjunction of preconditions (2) to (5).

Inspired by [6] and using the same notations, we propose in Table 3 the grammar of axiom <guarded reconfiguration> for *guarded reconfigurations*. Let <*ope*> represent a primitive reconfiguration operation, also called *primitive statement*. We extend the set of primitive reconfiguration operations with the *skip* operation, which does not induce any change on a given configuration. Hence, for any postcondition R , we have $wp(skip, R) = R$. Afterwards, like in [6], the semantics of the “;” operator is given by $wp(S_1; S_2, R) = wp(S_1, wp(S_2, R))$ where S_1 and S_2 are statements.

Guarded reconfiguration sets are used to define the alternative and the repetitive constructs; these sets are not statements. In a nutshell, the alternative construct selects for execution only guarded lists with a true guard, whereas, the repetitive construct selects for execution guarded lists with a true guard and is

Table 3: Guarded reconfigurations grammar

$\langle \text{guarded reconfiguration} \rangle$	$::= \langle \text{guard} \rangle \rightarrow \langle \text{guarded list} \rangle$
$\langle \text{guard} \rangle$	$::= \langle \text{boolean expression} \rangle$
$\langle \text{guarded list} \rangle$	$::= \langle \text{statement} \rangle \{ ; \langle \text{statement} \rangle \}$
$\langle \text{guarded reconfiguration set} \rangle$	$::= \langle \text{guarded reconfiguration} \rangle \{ \parallel \langle \text{guarded reconfiguration} \rangle \}$
$\langle \text{alternative construct} \rangle$	$::= \mathbf{if} \langle \text{guarded reconfiguration set} \rangle \mathbf{fi}$
$\langle \text{repetitive construct} \rangle$	$::= \mathbf{do} \langle \text{guarded reconfiguration set} \rangle \mathbf{od}$
$\langle \text{statement} \rangle$	$::= \langle \text{alternative construct} \rangle \mid \langle \text{repetitive construct} \rangle \mid \langle \text{ope} \rangle$

repeated until none of the guards is true. If a guarded reconfiguration set is made of more than one guarded reconfiguration, they are separated by the \parallel operator⁴.

To present the semantics of the alternative construct, let IF denote $\mathbf{if} B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n \mathbf{fi}$ and BB denote $(\exists i : 1 \leq i \leq n : B_i)$, then $wp(IF, R) = BB \wedge (\forall i : 1 \leq i \leq n : B_i \Rightarrow wp(S_i, R))$. For the repetitive construct, let DO denote $\mathbf{do} B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n \mathbf{do}$. Let $H_0(R) = R \wedge \neg BB$ and for $k > 0$, $H_k(R) = wp(IF, H_{k-1}(R)) \vee H_0(R)$, then $wp(DO, R) = \exists k : k \geq 0 : H_k(R)$. Intuitively, $H_k(R)$ is the weakest precondition guaranteeing termination after at most k selections of a guarded list, leaving the system in a configuration such that R holds. Let $\mathcal{R}_{run} = \mathcal{R} \cup \{run\}$ be a set of operations, where \mathcal{R} is a finite set of guarded reconfigurations instantiated wrt. the system under consideration, and run is the name of a generic action representing all the running operations⁵ of the component-based system.

Definition 3 (Reconfiguration model). *The operational semantics of a component-based system 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 configurations, $\mathcal{C}^0 \subseteq \mathcal{C}$ is a set of initial configurations, $\rightarrow \subseteq \mathcal{C} \times \mathcal{R}_{run} \times \mathcal{C}$ is the reconfiguration relation obeying $wp()$ predicates, and $l : \mathcal{C} \rightarrow CP$ is a total interpretation function.*

Let us note $c \xrightarrow{ope} c'$ for $(c, ope, c') \in \rightarrow$. Given the model $S = \langle \mathcal{C}, \mathcal{C}^0, \mathcal{R}_{run}, \rightarrow, l \rangle$, a path σ of S is a sequence of configurations c_0, c_1, c_2, \dots such that $\forall i \geq 0. \exists ope_i \in \mathcal{R}_{run}. (c_i \xrightarrow{ope_i} c_{i+1})$. An execution is a path σ in Σ s.t. $\sigma(0) \in \mathcal{C}^0$. We write $\sigma(i)$ to denote the i -th configuration of σ . The notation σ_i denotes the suffix path $\sigma(i), \sigma(i+1), \dots$, and σ_i^j denotes the segment path $\sigma(i), \sigma(i+1), \dots, \sigma(j-1), \sigma(j)$. Let Σ denote the set of paths, and Σ^f ($\subseteq \Sigma$) the set of finite paths. A configuration c' is reachable from c when there is a path $\sigma = c_0, c_1, \dots, c_n$ in Σ^f s.t. $c = c_0$ and $c' = c_n$ with $n \geq 0$. Let c be a configuration, the set of all configurations reachable from c is denoted $reach(c)$. This notion can be lifted from configurations to sets of configurations by $reach(\mathcal{C}) = \{reach(c) \mid c \in \mathcal{C}\}$.

Proposition 1 (Consistency propagation). *Given $\mathcal{C}^0 \subseteq \mathcal{C}$, $consistent(\mathcal{C}^0)$ implies $consistent(reach(\mathcal{C}^0))$.*

⁴ As in [6], the order in which guarded reconfigurations appear is semantically irrelevant.

⁵ The normal running of different components also changes the architecture, e.g., by modifying parameter values or stopping components.

Proof (sketch). We start the proof (see Appendix A for a more complete proof) by showing that each primitive operation ope preserves configuration consistency. We use this result to establish (by induction) that a guarded reconfiguration having a sequence of primitive statements in its guarded list also preserves consistency.

This allows us to show that guarded reconfigurations having a statement based on a guarded reconfiguration set made only of primitive statements ($G \rightarrow \mathbf{fi} \textit{grs} \mathbf{fi}$ or $G \rightarrow \mathbf{do} \textit{grs} \mathbf{od}$, where \textit{grs} denotes $B_0 \rightarrow ope_0 \parallel B_1 \rightarrow ope_1 \parallel \dots \parallel B_n \rightarrow ope_n$) also preserve consistency using only hypothesis on the statements' preconditions and postconditions.

Therefore, with the same reasoning, considering general (i.e., primitive or non primitive) statements instead of only primitive ones and using only hypothesis on statements' preconditions and postconditions, we can prove that consistency is preserved *a)* for guarded reconfigurations having a guarded list composed of a sequence of (non primitive) statements ($G \rightarrow S_0; S_1; \dots; S_n$) and *b)* for guarded reconfigurations having as guarded list a statement ($G \rightarrow \mathbf{fi} \textit{grs} \mathbf{fi}$ or $G \rightarrow \mathbf{do} \textit{grs} \mathbf{od}$, where \textit{grs} denotes $B_0 \rightarrow S_0 \parallel B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n$). \square

4 Interpreted Architecture Model

In the specification model, primitive operations and guarded reconfigurations were left abstract enough and *run* was uninterpreted. A formal semantics for the component-based system with interpreted operations can be obtained by enriching the configurations with more precise memory states and the effect of these actions upon memory.

4.1 Interpreted Configurations and Reconfigurations

Let us consider a set (infinite, in general) $GM = \{u, \dots\}$ of shared global memory states, and a set (infinite, in general) $LM = \{v, \dots\}$ of memory states local to a given component. These memory states are read and modified by the primitive and non-primitive reconfigurations, and also by actions implementing *run*.

Interpreted configurations. In addition to already interpreted parameters and interfaces (cf. [5] for more detail), the state of components can be described more precisely by using local memory states. The set of the interpreted states of components is the least set $State_{\mathcal{I}}$ s.t. if s_1, \dots, s_n are elements in $State_{\mathcal{I}}$ ⁶, $v_1, \dots, v_n \in LM$ are local memory states, then $((s_1, v_1), \dots, (s_n, v_n))$ is in $State_{\mathcal{I}}$. Then, the set of the interpreted configurations $\mathcal{C}_{\mathcal{I}}$ is defined by $GM \times State_{\mathcal{I}}$.

Interpreted transitions. Our basic assumption is that all primitive actions have a deterministic effect upon the local and global memory, always terminate (either normally or exceptionally), and are effective. For each primitive reconfiguration operation ope , the corresponding interpreted reconfiguration, denoted by \overline{ope} , has equivalent or stronger preconditions, such that all constructs behave deterministically. A non-deterministic global behaviour is produced by the arbitrary interleaving of components.

⁶ Viewed as a relation.

Formally, all the actions $ope \in \mathcal{R}_{run}$ are interpreted as mappings \overline{ope} from $GM \times LM$ into itself. Additionally, there are some actions specific to the interpretation, \mathcal{R}_{int} , for example for testing guards. We say that $\mathcal{I} = (GM, LM, (\overline{ope})_{ope \in \mathcal{R}_{run} \cup \mathcal{R}_{int}})$ is an interpretation of the underlying \mathcal{R}_{run} . Let $\mathcal{I}_{\mathcal{R}_{run}}$ denote the class of all interpretations. This construction leads to

Definition 4 (Interpreted reconfiguration model). *The interpreted operational semantics of component-based system is defined by the labelled transition system $S_{\mathcal{I}} = \langle \mathcal{C}_{\mathcal{I}}, \mathcal{C}_{\mathcal{I}}^0, \mathcal{R}_{run_{\mathcal{I}}}, \rightarrow_{\mathcal{I}}, l_{\mathcal{I}} \rangle$ where $\mathcal{C}_{\mathcal{I}}$ is a set of configurations together with their memory states, $\mathcal{C}_{\mathcal{I}}^0$ is a set of initial configurations, $\mathcal{R}_{run_{\mathcal{I}}} = \{\overline{ope} \mid ope \in \mathcal{R}_{run} \cup \mathcal{R}_{int}\}$, $\rightarrow_{\mathcal{I}} \subseteq \mathcal{C}_{\mathcal{I}} \times \mathcal{R}_{run_{\mathcal{I}}} \times \mathcal{C}_{\mathcal{I}}$ is the interpreted reconfiguration relation, and $l_{\mathcal{I}} : \mathcal{C}_{\mathcal{I}} \rightarrow CP$ is a total interpretation function.*

It is easy to see that, by construction, $\text{consistent}(\mathcal{C}_{\mathcal{I}}^0)$. Moreover, if $\text{consistent}(c)$ and $c \xrightarrow{\overline{ope}}_{\mathcal{I}} c'$ then $\text{consistent}(c')$.

4.2 Compatible Interpretation

To establish links between the reconfiguration model and the corresponding interpreted model, we propose to use a version of the classical τ -simulation quasi-ordering [12], while relabeling the operations in \mathcal{R}_{int} by τ . For all $ope \in \mathcal{R} \cup \{\epsilon\}$, where ϵ denotes the empty word, we write $c \xrightarrow{ope} c'$ when there are $n, m \geq 0$ such that $c \xrightarrow{\tau^n ope \tau^m} c'$.

Definition 5 (d -simulation). *Let $S_1 = \langle \mathcal{C}_1, \mathcal{C}_1^0, \dots \rangle$ and $S_2 = \langle \mathcal{C}_2, \mathcal{C}_2^0, \dots \rangle$ be two models over \mathcal{R} . A binary relation $\sqsubseteq_d \subseteq \mathcal{C}_1 \times \mathcal{C}_2$ is a d -simulation iff, for all ope in $\mathcal{R} \cup \{\epsilon\}$, $(c_1, c_2) \in \sqsubseteq_{\tau}$ implies 1) whenever $c_1 \xrightarrow{ope}_1 c'_1$, then there exists $c'_2 \in \mathcal{C}_2$ such that $c_2 \xrightarrow{ope}_2 c'_2$ and $(c'_1, c'_2) \in \sqsubseteq_d$, and 2) $c_1 \not\xrightarrow{ope}$ implies $c_2 \not\xrightarrow{ope}$.*

We write $S_1 \sqsubseteq_d S_2$ when $\forall c_1^0 \in \mathcal{C}_1^0 \exists c_2^0 \in \mathcal{C}_2^0. (c_1^0, c_2^0) \in \sqsubseteq_d$.

Let us consider interpreted reconfiguration operations in $\mathcal{R}_{run_{\mathcal{I}}}$ and the corresponding non-interpreted counterpart in \mathcal{R}_{run} . When relabelling the operations in \mathcal{R}_{int} by τ , we can state—modulo the overline notation—that the more abstract model τ -simulates the interpreted model (because of the non-determinism when testing guards in the non-interpreted model); nevertheless, this simulation respects non-divergency.

Theorem 1 (Compatibility). $S_{\mathcal{I}} \sqsubseteq_d S$.

Proof (sketch). There are two cases for $ope \in \mathcal{R}_{run} \cup \mathcal{R}_{int}$. As τ 's covering operations in \mathcal{R}_{int} are introduced to evaluate guards of sequences of guarded reconfigurations, they do not form infinite cycles of τ -transitions. So, there always must be a way out of these cycles, if any, by a transition of label \overline{ope} .

By construction any primitive reconfiguration operation of the interpreted model has preconditions equivalent to or stronger than its counterpart in the non-interpreted model. This way, by using hypothesis on weakest preconditions in [6], we can prove that guarded reconfigurations composed of primitive statements,

$G \rightarrow \bar{s}$, with $\bar{s} \in \mathcal{R}_{run_{\mathcal{I}}} \setminus \mathcal{R}_{int}$ have preconditions equivalent to or stronger than the corresponding statement $s \in \mathcal{R}_{run}$. Consequently, starting from initial configurations, for any $c_1 \in \mathcal{C}_{\mathcal{I}}$, if $\mathbf{consistent}(c_1)$ there is $c_2 \in \mathcal{C}$ s.t. $\mathbf{consistent}(c_2)$, and if a guarded reconfiguration $G \rightarrow \bar{s}$ is applied to c_1 there exists a guard G' , s.t. $G \Rightarrow G'$ and $G' \rightarrow s$ applies to c_2 . Moreover, the consistent target configurations are in \sqsubseteq_d too because of their guards.

If no *ope* can be performed in $c_1 \in \mathcal{C}_{\mathcal{I}}$ after having tested some guards covered by τ , c_1 is not consistent, and consequently neither is $c_2 \in \mathcal{C}$. At this step, only several primitive reconfigurations can be applied, as their preconditions are equivalent, no *ope* can be performed in c_2 either. \square

4.3 Property Preservation

Theorem 1 can be exploited for property preservation. For example, as the reachability properties are compatible with \sqsubseteq_d , this leads us, consequently, to:

Proposition 2. *If configuration c is not reachable in S , it is not reachable in any $S_{\mathcal{I}}$. Conversely, if configuration c is reachable in S , there exists an interpretation \mathcal{I} such that c is reachable in $S_{\mathcal{I}}$.*

In addition, safety properties expressed via non-reachability properties can be ensured. Moreover, as a consequence of Theorem 1 and Propositions 1 and 2, we can state:

Proposition 3. *Let $S_{\mathcal{I}} = \langle \mathcal{C}_{\mathcal{I}}, \mathcal{C}_{\mathcal{I}}^0, \mathcal{R}_{run_{\mathcal{I}}}, \rightarrow_{\mathcal{I}}, l_{\mathcal{I}} \rangle$ be the interpreted model and $S = \langle \mathcal{C}, \mathcal{C}^0, \mathcal{R}_{run}, \rightarrow, l \rangle$ the specification model. Given $\mathcal{C}_{\mathcal{I}}^0 \subseteq \mathcal{C}_{\mathcal{I}}$, if $S_{\mathcal{I}} \sqsubseteq_d S$ then $\mathbf{consistent}(\mathcal{C}_{\mathcal{I}}^0)$ implies $\mathbf{consistent}(\mathit{reach}(\mathcal{C}_{\mathcal{I}}^0))$.*

It must be noticed that differently from [3], we do not assume that the reconfigurations always make evolve the component assembly from one consistent architecture to another consistent architecture, only through a path of consistent configurations. Indeed, this assumption seems to be too strong notably wrt. primitive reconfigurations.

5 Implementation and Architecture Conformance

5.1 Implementation Protocol

We developed a prototype tool, contained in a java package named *cbdr*⁷, supporting the interpreted reconfiguration model to design and simulate component-based systems with dynamic reconfigurations. Using generic java classes, we can use our implementation to perform reconfigurations on applications deployed using Fractal [7] or FraSCAti [8]. The Fractal framework is based on a hierarchical and reflective component model. Its goal is to reduce the development, deployment, and maintenance costs of software systems in general⁸. FraSCAti is an open-source

⁷ *cbdr* stands for Component-Based System Dynamic Reconfiguration

⁸ <http://fractal.ow2.org/tutorial/index.html>

implementation of the *Service Component Architecture*⁹ (SCA). It can be seen as a framework having a Fractal base with an extra layer implementing SCA specifications. In [8], a smart home scenario illustrates the capabilities and the various reconfigurations of the FraSCAti platform.

Figure 3 shows the *cbstr* interface displaying a given state of the VM from our running example developed using Fractal (top frame). The left frame shows the various states of the run under scrutiny, whereas the bottom frame can be used to display various information such as the evolution of parameters of the model, console output, or the outcome of reconfigurations performed.

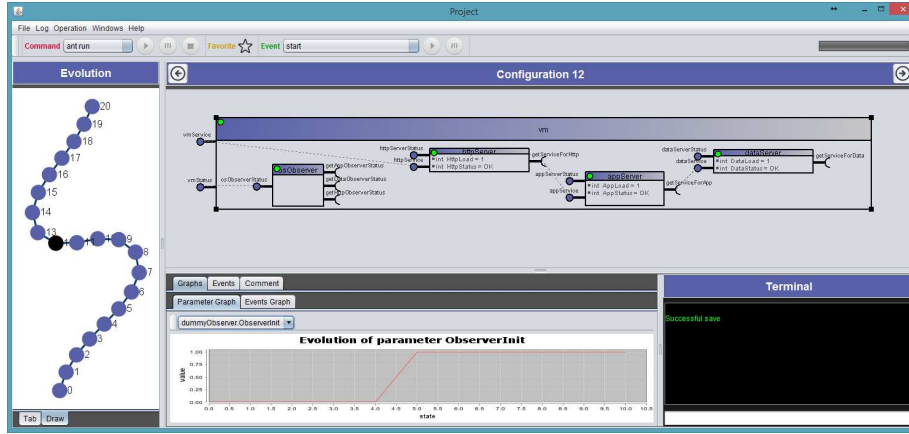


Fig. 3: Model of the VM component-based system displayed in our interface

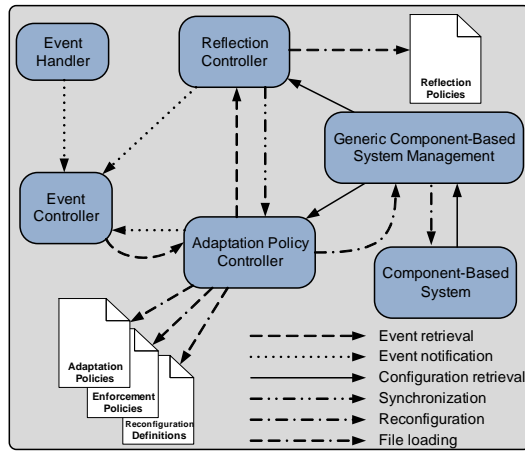


Fig. 4: *cbstr* Implementation Architecture

This interface allows the monitoring of a component-based system and the generation of (external) events during a run of *cbstr*, but can also be used to analyse the logs of a run already performed. It is interesting to note that primitive, as well as, non primitive reconfiguration operations can be performed and analysed.

Thanks to this application, in addition to the above-mentioned functionalities, we are able to perform adaptations using dynamic reconfig-

⁹ <http://www.oasis-openca.org/sca>

urations triggered by temporal properties at runtime, as described in [2]. In this case, the implementation (see Fig. 4) works as follows: *a*) adaptation policies are loaded and applied using a control loop, *b*) FTPL¹⁰ expressions are evaluated and (if any) candidate reconfigurations are ordered by priority using fuzzy logic values embedded in adaptation policies, *c*) candidate reconfigurations are applied to the component-based system model using our reconfiguration semantics in order to verify that the corresponding target configuration does not violate any of the properties to enforce, and *d*) the target configuration obtained using the reconfiguration with highest priority that does not violate any of the properties to enforce is applied to the component based system using a protocol similar to the one described in [3]. The fact that we are using temporal properties based on architectural relations as well as internal and external events allows us to significantly reduce communication overhead *a*) by, as in [13], using decentralised evaluation of temporal properties or *b*) by allowing the user to submit simultaneous (external) events to the system, as explained below.

5.2 Architecture Conformance

The reconfiguration model is a correct approximation of the more realistic interpreted implementation model. This fact can be expressed by using the notion of conformance of the component architecture model. Basically, following the most commonly used *ioco* relation in [14], an implementation $S_{\mathcal{I}}$ is conformant to its specification S if, after a trace of S , one should foresee the output of $S_{\mathcal{I}}$ in S , and the implementation is authorised to reach a state where it cannot produce any output only if this is the case in the specification too.

Using various simulation relations permits expressing trace-inclusion-based conformance and stronger conformance relations at the level of transition systems. Thus, thanks to the proof arguments of Theorem 1, and the subsequent trace inclusion modulo τ , we have the following conformance result, with $S_{\mathcal{I}_{cbsdr}}$ being the *cbsdr* implementation.

Proposition 4. $S_{\mathcal{I}_{cbsdr}}$ is conformant to S .

5.3 Running Example

We consider a VM represented, as in Fig. 1, as a composite component *virtualMachine* that may contain sub-components representing services *httpServer*, *appServer*, or *dataServer* of an application. This VM may also contain *observers* that are sub-components used to monitor services. The sub-component *osObs* is used to monitor the Operating System of the VM and can be bound to the sub-components *httpObs*, *appObs*, or *dataObs* used respectively to monitor the services of the *httpServer*, *appServer*, and *dataServer* sub-components.

¹⁰ 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.

The Fractal and FraSCAti versions of the VM example can be controlled by our implementation using external events as *init*, *manage*, *setdata*, etc., to (respectively) initialise the VM, monitor the VM, or set the data server of the VM up. If the VM is monitored, it is described as *managed*, otherwise it is said to be *unmanaged*. Depending of the service to provide and the state of the VM (managed vs. unmanaged), only the necessary component should be added.

For example, let us consider a managed VM providing only the HTTP service: it contains the *httpServer* component and, since it is managed, it also contains the *osObs* and the *httpObs* components. Therefore, the generation of the *setdata* external event triggers (via adaptation policies) the addition of the *dataServer* and *dataObs* components. Of course, if the initial VM was unmanaged, the generation of the *setdata* external event would only result in the addition of the *dataServer* component. Nevertheless, in this case (unmanaged HTTP VM), the generation of the *setdata* and *manage* external events would result in a VM containing all the components pertaining to a managed VM providing the HTTP and the DATA services (i.e, *httpServer*, *dataServer*, *osObs*, *httpObs*, and *dataObs*).

This is due to the fact that we use FTPL temporal logic expressions as “**after** *unsetdata* ((**always** \top) **until** *setdata*)” to guarantee that, in case of opposite events like *setdata* and *unsetdata*, the corresponding expression is potentially true until the occurrence of the opposite events. This way, the ordered sequence of events *init*, *manage*, *sethttp*, *setapp*, and *setdata* is equivalent to a single communication containing all these events at once; this significantly reduces communication overhead.

5.4 Other Examples

To illustrate how the *csdr* tool works, we present below two examples: a small http server, and a model of the location component of the *cycab*, an autonomous vehicle. This latter example confirms that not only pertinent reconfigurations can be triggered, but also reconfigurations leading to unwanted behaviours are avoided. Finally, we conclude this section by presenting some results about the CPU consumption of the *csdr* tool used with both Fractal and FraSCAti frameworks.

Http Server. Figure 5 shows an experimentation with the http server composite component during which, as in [15], http requests were simulated. Depending on the *load* and *request deviation* to measure whether or not requests are similar, it may make sense to add a cache (the need can be *low*, *medium*, or *high* determining the size of the cache), and an additional file server.

Interestingly, response times measured when our http server is controlled and adapted by the *csdr* application match almost exactly the times measured (under similar load and request deviation patterns) for a http server having a cache (of size *high*) and two file servers. No memory nor disk overhead were noted.




Fig. 5: Experiment with the http server composite component

Cycab. Figure 6 uses the model of the location system of an autonomous car. Thanks to adaptation using temporal properties at runtime, we can remove the *gps* or *wifi* location components to save energy when needed (e.g., the *gps* component does not work in tunnels — between *entry* and *exit*).




Fig. 6: Experiment with the cycab location composite component

The run represented in Fig. 6 shows a consumption of energy around 32% lower using adaptation (empty dashed red graph) compared to a run not using it (full dashed blue graph). It is important to notice that when the vehicle is in a tunnel, the *cbstr* tool prevents the occurrence of the reconfiguration that would normally add the *gps* component when the power level is high. The reader interested in a more detailed description is referred to [2].

CPU Overhead. We tested our implementation on the above-mentioned examples using both Fractal and FraSCAti framework. More than 300 tests were performed to assess the resources overhead caused by our implementation. Table 4 summarises the increase of

Table 4: Measured increase of CPU usage expressed in percent ($\bar{\mu} \pm \sigma$)

Framework	Fractal	FraSCAti
CPU User time	17 \pm 3	11 \pm 2
CPU System time	2 \pm 2	14 \pm 2
Percent of CPU	17 \pm 2	15 \pm 7

CPU usage when adaptation is used compared to similar runs not using any adaptation mechanism. CPU overhead is expressed in Table 4 in the format $\bar{\mu} \pm \sigma$ with $\bar{\mu}$ being the average and σ the standard deviation.

6 Related Work and Conclusion

6.1 Related Work

Self-adaptation is an important and active research field with applications in various domains [1]. This roadmap emphasises an important challenge consisting in bridging the gap between the design and the implementation of self-adaptive systems. In [2] component-based systems reconfiguration was performed at runtime using adaptation policies triggered by temporal patterns. The reconfigurations considered, however, were merely sequences of primitive reconfiguration operations. In the present paper, since we use the alternative and the repetitive constructs to compose reconfigurations, a given reconfiguration can have different outcomes, depending on the context, or due to non-deterministic mechanisms. It is not only a static sequence of reconfiguration instructions (as it is the case in [2,7,8,16]), but a truly *dynamic* reconfiguration. Differently from [3], we do not assume that the reconfigurations always lead the component assembly to evolve from one consistent architecture to another consistent architecture.

Version consistency was introduced in [16] to minimise the interruption of service (*disruption*) and the delay with which component-based (distributed) systems are updated (*timeliness*) by means of reconfigurations. It qualifies a state where transactions within the system are such that a given reconfiguration may not disrupt the system and occur in bounded time; version consistency was inspired by *quiescence* [17] and *tranquility* [18] with the intent to gather the best of both notions. Unlike [16,17,18], we only consider architectural constraints as preconditions to apply guarded reconfigurations; this way, by considering components as black boxes, the separation of concerns principle is respected. The applicative consistency (related to transactions within the system or external events) can be maintained at runtime using adaptation policies mechanisms as in [2] for centralised system and in [13] for decentralised or distributed systems.

Following [19], our notion of consistency can be viewed as a specific architecture style. Nevertheless, when using graph grammars, we represent interfaces types of the component-based systems by specific graph nodes, this way, like in [20], we can monitor (temporal) properties at the interface level.

Let us remark that the present work is motivated by other frameworks that support the development of components. For example, experimenting with our VM example within the *GROOVE* environment [21] leads us to the presentation of paths with transitions labelled by the primitive reconfiguration operations being performed. Consequently, consistency and conformity issues are pertinent to *GROOVE* too.

6.2 Conclusion

Inspired by [6], we proposed a grammar for guarded reconfigurations. This allowed us to build reconfigurations based on primitive reconfiguration operations using sequences of reconfigurations as well as the alternative and the repetitive constructs. The ability to determine weakest preconditions for the application of reconfigurations enabled us to prove that these guarded reconfigurations preserve configuration consistency.

This way, a conformance relation can be established to validate implementations of component-based systems architectural models using either our java-based *cbstr* application or the *GROOVE* graph transformation tool. This makes these tools applicable to build some parts of state space of reachable graphs, i.e., configurations, and thereby derive information about the system. Furthermore, one of the key advantages of this work is that it is readily applicable to practical reconfiguration operations.

As a future work, we intend to exploit the results of the present paper to extend adaptation policies defined in [2] with guarded reconfigurations. Then, we could aim to perform sound and complete compositions of such adaptation policies. This would permit us to move further toward our overall goal, which is the adaptation of component-based system at runtime using adaptation policies based on temporal logic properties.

References

1. De Lemos, R., Giese, H., Müller, H.A., Shaw, M., Andersson, J., Litoiu, M., Schmerl, B., Tamura, G., Villegas, N.M., Vogel, T., et al.: Software engineering for self-adaptive systems: A second research roadmap. In: *Software Engineering for Self-Adaptive Systems II*. Springer (2013) 1–32
2. 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.: *10th Int. Symp. on Formal Aspects of Component Software, FACS 2013, Nanchang, China, October 27-29, 2013, Revised Selected Papers*. Volume 8348 of LNCS. Springer (2014) 234–253
3. Boyer, F., Gruber, O., Pous, D.: Robust reconfigurations of component assemblies. In: *Int. Conf. on Software Engineering, ICSE '13, Piscataway, NJ, USA, IEEE Press (2013)* 13–22
4. Myllärniemi, V., Ylikangas, M., Raatikainen, M., Pääkkö, J., Männistö, T., Aaltonen, T.: Configurator-as-a-service: tool support for deriving software architectures at runtime. In: *The WICSA/ECSCA 2012 Companion Volume, ACM (2012)* 151–158
5. Lanoix, A., Dormoy, J., Kouchnarenko, O.: Combining proof and model-checking to validate reconfigurable architectures. *ENTCS* **279** (2011) 43–57
6. Dijkstra, E.W.: Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM* **18** (1975) 453–457
7. Bruneton, E., Coupaye, T., Leclercq, M., Quéma, V., Stefani, J.B.: The fractal component model and its support in java. *Software: Practice and Experience* **36** (2006) 1257–1284

8. 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
9. Dormoy, J., Kouchnarenko, O., Lanoix, A.: Using temporal logic for dynamic reconfigurations of components. In Barbosa, L., Lumpe, M., eds.: 7th Int. Wshop on Formal Aspects of Component Software, FACS 2010, Guimarães, Portugal, October 14-16, 2010, Revised Selected Papers. Volume 6921 of LNCS. Springer Berlin Heidelberg (2012) 200–217
10. Hamilton, A.G.: *Logic for mathematicians*. Cambridge University Press (1988)
11. Hoare, C.A.R.: An axiomatic basis for computer programming. *Communications of the ACM* **12** (1969) 576–580
12. Milner, R.: *Communication and concurrency*. Prentice-Hall, Inc. (1989)
13. Kouchnarenko, O., Weber, J.F.: Decentralised evaluation of temporal patterns over component-based systems at runtime. In Lanese, I., Madelaine, E., eds.: 11th Int. Symp. on Formal Aspects of Component Software, FACS 2014, Bertinoro, Italy, September 10-12, 2014, Revised Selected Papers. Volume 8997 of LNCS. Springer International Publishing (2015) 108–126
14. Tretmans, J.: Test generation with inputs, outputs and repetitive quiescence. *Software - Concepts and Tools* **17** (1996) 103–120
15. Chauvel, F., Barais, O., Plouzeau, N., Borne, I., Jézéquel, J.M.: Composition et expression qualitative de politiques d’adaptation pour les composants Fractal. In: Actes des Journées nationales du GDR GPL 2009. (2009)
16. Ma, X., Baresi, L., Ghezzi, C., Panzica La Manna, V., Lu, J.: Version-consistent dynamic reconfiguration of component-based distributed systems. In: The 19th ACM SIGSOFT Symp. and the 13th European Conf. on Foundations of software engineering, ACM (2011) 245–255
17. Kramer, J., Magee, J.: The evolving philosophers problem: Dynamic change management. *Software Engineering, IEEE Transactions on* **16** (1990) 1293–1306
18. Vandewoude, Y., Ebraert, P., Berbers, Y., D’Hondt, T.: Tranquility: A low disruptive alternative to quiescence for ensuring safe dynamic updates. *Software Engineering, IEEE Transactions on* **33** (2007) 856–868
19. Le Metayer, D.: Describing software architecture styles using graph grammars. *IEEE Transactions on Software Engineering* **24** (1998) 521–533
20. Kähkönen, K., Lampinen, J., Heljanko, K., Niemelä, I.: The lime interface specification language and runtime monitoring tool. In Bensalem, S., Peled, D.A., eds.: Int. Wshop on Runtime Verification, RV’09. Volume 5779 of LNCS. Springer Berlin Heidelberg (2009) 93–100
21. 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

A Proof of Proposition 1

We start the proof by establishing that each primitive operation ope preserves configuration consistency. This means, for R being a postcondition of ope , that we have $CC \wedge wp(ope, R) = wp(ope, CC \wedge R)$. We show this result for $add(,)$, the proof is similar for the other primitive operations. Let c be such that $consistent(c)$ and the preconditions of $add(,)$ hold on c : Then, the transition $c \xrightarrow{add} c'$ leads to configuration c' such that $consistent(c')$, i.e., that the postconditions of $add(,)$ satisfy the consistency constraints of Table 1 too; formally, $(l(c) \Rightarrow CC \wedge wp(add, R_{add})) \wedge (c \xrightarrow{add} c') \Rightarrow (l(c') \Rightarrow CC \wedge R_{add})$. Indeed, as the *Parent* relation from the postcondition (1) is not involved in **(CC.1)**, **(CC.4)** to **(CC.9)**, **(CC.11)**, and **(CC.12)**, these constraints hold on c' too. For the remaining constraints, one has:

- (CC.2)**: As precondition (5) of Table 2 ensures that the parent component $comp_2$ has no parameters, **(CC.2)** holds on c' with $(comp_1, comp_2)$ added to *Parent* (cf. (1));
- (CC.3)**: Precondition (4) of Table 2 means that $comp_2$ cannot be a descendant of $comp_1$, thus preventing a cycle in the *Parent* relation for c' when $comp_2$ becomes a parent of $comp_1$;
- (CC.10)**: There are two cases: Either there already was a delegation relation between interfaces of $comp_1$ and $comp_2$ on c before the application of the $add(,)$ operation, or not. In the latter case the constraint **(CC.10)** trivially holds on c' . In the former case, since $consistent(c)$, the *Parent* relation already had $(comp_1, comp_2)$ with well-typed interfaces for c , and the application of $add(,)$ does not change the types and the relation, therefore the constraint holds on c' .

Let be $c \in reach(\mathcal{C}^0)$; by definition, there exists $c_0 \in \mathcal{C}^0$ and a sequence of operations from \mathcal{R}_{run} to ultimately reach c . By definition, there also exists a sequence of primitive operations $ope_0, ope_1, \dots, ope_{n-1}$ and a set of intermediate configurations $\mathcal{C}' = \{c_1, c_2, \dots, c_{n-1}\}$ ¹¹ such that $c_0 \xrightarrow{ope_0} c_1, c_1 \xrightarrow{ope_1} c_2, \dots, c_{n-1} \xrightarrow{ope_{n-1}} c$, where, for $0 \leq i \leq n-1$, c_i (resp. c_{i+1}) meets the preconditions (resp. postconditions) of ope_i (c_n standing for c). Indeed, if this sequence of primitive operations or \mathcal{C}' would not exist, c would not be reachable from any configuration in \mathcal{C}^0 .

Now, let us prove that a guarded reconfiguration having a sequence of primitive statements in its guarded list preserves consistency. Let gl_n be a guarded list composed of $n \geq 0$ primitive operations, i.e., $gl_n = ope_0; ope_1; \dots; ope_n$, with R_i and R_{i+1} being respectively preconditions and postconditions of ope_i , we note $CC_i = CC \wedge R_i$. Let us prove by induction on n that $CC_0 = wp(gl_n, CC_{n+1})$. For $n = 0$, we have $gl_n = ope_0$ and $CC_0 = wp(gl_0, CC_1)$. Let us now consider

¹¹ Note that \mathcal{C}' is not necessarily a subset of \mathcal{C} . For example, if each operation of \mathcal{R} is a sequence of two primitive operations, the intermediary configuration with odd index, i.e., c_1, c_3, \dots , would not belong to \mathcal{C} and $\mathcal{C}' \not\subseteq \mathcal{C}$.

$gl_{n+1} = gl_n; ope_{n+1}$; we have $wp(gl_{n+1}, CC_{n+2}) = wp(gl_n, wp(ope_{n+1}, CC_{n+2}))$. Since $CC_0 = wp(gl_n, CC_{n+1})$ and $CC_{n+1} = wp(ope_{n+1}, CC_{n+2})$, we have, by definition [6], $CC_0 = wp(gl_n, CC_{n+1}) = wp(gl_n, wp(ope_{n+1}, CC_{n+2}))$.

We prove, below, that guarded reconfigurations having a non primitive statement based on a guarded reconfiguration set made only of primitive statements ($G \rightarrow \mathbf{fi}$ *grs* \mathbf{fi} or $G \rightarrow \mathbf{do}$ *grs* \mathbf{od} , where *grs* denotes $B_0 \rightarrow ope_0 \parallel B_1 \rightarrow ope_1 \parallel \dots \parallel B_n \rightarrow ope_n$) also preserve consistency using only hypothesis on the statements' preconditions and postconditions.

Let us consider a guarded reconfiguration set *grs* based on guarded reconfigurations containing guarded lists made only of primitive statements. This reconfiguration set, *grs* denotes $B'_1 \rightarrow S_1 \parallel \dots \parallel B'_n \rightarrow S_n$, with B'_i being a boolean and $S_i = ope_0^i; ope_1^i; \dots; ope_{n_i}^i$, where n_i represent the number of primitive statements ($ope_0^i, ope_1^i, \dots, ope_{n_i}^i$) of the guarded list S_i , and R_j^i (resp. R_{j+1}^i) represents the precondition (postcondition) of ope_j^i , for $0 < i \leq n$ and $0 \leq j \leq n_i$ (resp. $0 < j \leq n_i + 1$).

Since R_0^i is the precondition of S_i and we suppose the configuration before the application of S_i to be consistent, we rewrite *grs* as $B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n$, with $B_i = B'_i \wedge CC \wedge R_0^i$. We also define $BB = (\exists i : 1 \leq i \leq n : B_i)$, as well as, the sets $I = \{i \in \mathbb{N}. 1 \leq i \leq n\}$ and $I_\top = \{i \in I. B_i\}$.

Alternative Construct Let *IF* denote $\mathbf{if} B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n \mathbf{fi}$.

By definition, $wp(IF, R) = BB \wedge \forall i \in I : B_i \Rightarrow wp(S_i, R)$. We established before that, for S_i being a sequence of primitive statement, $CC \wedge R_0^i \Rightarrow wp(S_i, CC \wedge R_{n_i+1}^i)$; then, by definition $B_i \Rightarrow wp(S_i, CC \wedge R_{n_i+1}^i) \Rightarrow wp(S_i, CC)$. This means that $wp(IF, CC) = BB \wedge \forall i \in I : B_i \Rightarrow wp(S_i, CC)$, which can be enough to prove that consistency is preserves by the alternative construct.

It is possible, however, to establish a stronger postcondition, $CC \wedge \bigwedge_{i \in I_\top} R_{n_i+1}^i$, for the alternative construct by considering that each term of the conjunction $\bigwedge_{i \in I_\top} R_{n_i+1}^i$ is part of the postcondition of a guarded list eligible for execution.

Then, $wp(IF, CC \wedge \bigwedge_{j \in I_\top} R_{n_j+1}^j) = BB \wedge \forall i \in I : B_i \Rightarrow wp(S_i, CC \wedge R_{n_i+1}^i)$ because, by definition, $\forall i \in I_\top, B_i = \top$.

Therefore:

$$\begin{aligned} BB \wedge CC \wedge \bigwedge_{i \in I_\top} R_0^i &\Rightarrow BB \wedge \bigwedge_{i \in I_\top} wp(S_i, CC \wedge R_{n_i+1}^i) \\ &\Rightarrow BB \wedge (\forall i \in I : B_i \Rightarrow wp(S_i, CC \wedge R_{n_i+1}^i)) \\ &\Rightarrow wp(IF, CC \wedge \bigwedge_{j \in I_\top} R_{n_j+1}^j) \end{aligned}$$

As an example, we can denote by $\mathbf{if} B \mathbf{then} S \mathbf{fi}$ a particular case, written $\mathbf{if} B \rightarrow S \parallel \neg B \rightarrow skip \mathbf{fi}$, of the alternative construct which weakest precondition is $wp(\mathbf{if} B \mathbf{then} S \mathbf{fi}, CC \wedge (B \Rightarrow post_S)) = B \wedge wp(S, CC \wedge R_S)$, where R_S and $post_S$ are, respectively, the precondition and postcondition of S .

Repetitive Construct Let DO denote $\mathbf{do} B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n \mathbf{od}$. Let be $H_0(R) = R \wedge \neg B$, and for $k > 0$, let be $H_k(R) = wp(IF, H_{k-1}(R)) \vee H_0(R)$, where IF denotes the same guarded configuration enclosed by “**if fi**”. Then, by definition, we have $wp(DO, R) = \exists k : k > 0 : H_k(R)$.

This means that the weakest precondition of this construct guarantees proper termination after at most k selections of a guarded list, leaving the system in a state satisfying R . Let us consider l_0, l_1, \dots, l_k , such that, for $0 \leq j \leq k$, $1 \leq l_j \leq n$ and $S_{l_0}; S_{l_1}; \dots; S_{l_k}$, as the ordered sequence of statements selected during the duration of the construct until its termination. We proved before that such a sequence preserve consistency. Therefore $CC \wedge R_{n_{l_k+1}}^{l_k}$ is a valid postcondition and, since $CC \wedge R_{n_{l_k+1}}^{l_k} \Rightarrow CC \wedge \bigvee_{i \in I} R_{n_{i+1}}^i$, we have $wp(DO, CC \wedge R_{n_{l_k+1}}^{l_k}) \Rightarrow wp(DO, CC \wedge \bigvee_{i \in I} R_{n_{i+1}}^i)$.

We established before that, for S_i being a sequence of primitive statement, $CC \wedge R_0^i \Rightarrow wp(S_i, CC \wedge R_{n_{i+1}}^i)$; then $CC \wedge R_0^{l_0} \Rightarrow wp(S_{l_0}; S_{l_1}; \dots; S_{l_k}, CC \wedge R_{n_{l_k+1}}^{l_k})$.

Therefore:

$$\begin{aligned} CC \wedge \bigwedge_{i \in I} R_0^i &\Rightarrow CC \wedge R_0^{l_0} \\ &\Rightarrow wp(S_{l_0}; S_{l_1}; \dots; S_{l_k}, CC \wedge R_{n_{l_k+1}}^{l_k}) \\ &\quad \text{for any valid sequence } S_{l_0}; S_{l_1}; \dots; S_{l_k} \\ &\Rightarrow wp(DO, CC \wedge R_{n_{l_k+1}}^{l_k}) \\ &\Rightarrow wp(DO, CC \wedge \bigvee_{i \in I} R_{n_{i+1}}^i) \end{aligned}$$

This proves that the repetitive construct, applied to a guarded reconfiguration set based on guarded reconfigurations containing guarded lists made only of primitive statements, preserves consistency. It also provides stronger preconditions and postconditions that are used to conclude the full proof of Prop 1.

Conclusion Therefore, with the same reasoning, considering non primitive statements instead of primitive ones and using only hypothesis on statements' preconditions and postconditions, we can prove that consistency is preserved a) for guarded reconfigurations having a guarded list composed of a sequence of (non primitive) statements $(G \rightarrow S_0; S_1; \dots; S_n)$ and b) for guarded reconfigurations having as guarded list a statement $(G \rightarrow \mathbf{fi} \text{ grs } \mathbf{fi}$ or $G \rightarrow \mathbf{do} \text{ grs } \mathbf{od}$, where grs denotes $B_0 \rightarrow S_0 \parallel B_1 \rightarrow S_1 \parallel \dots \parallel B_n \rightarrow S_n$). \square