

Detecting Architectural Mismatches in Process Algebraic Descriptions of Software Systems

Marco Bernardo
Università di Torino
Dipartimento di Informatica
Corso Svizzera 185, 10149 Torino, Italy
bernardo@di.unito.it

Paolo Ciancarini, Lorenzo Donatiello
Università di Bologna
Dipartimento di Scienze dell'Informazione
Mura Anteo Zamboni 7, 40127 Bologna, Italy
cianca, donat@cs.unibo.it

Abstract

Formalizing the description of software systems helps to detect the presence of architectural mismatches that can arise when assembling software components together. In this paper we identify three causes of architectural mismatches: incompatibility between two components due to a single interaction, incompatibility between two components due to the combination of several interactions, and lack of interoperability among a set of components forming a cyclic topology. We then show how to deal with all of them within a uniform, process algebraic framework. We begin with the first two causes by strengthening a previously defined architectural compatibility check based on observational equivalences, in order to achieve a deadlock freedom result for the set of components interacting via a certain connection. We subsequently concentrate on the third cause by defining a novel architectural interoperability check based on observational equivalences, which guarantees absence of deadlock within a set of interacting components forming a cyclic topology. We finally assess the adequacy of our architectural interoperability check by applying it to the description of a cruise control system.

1. Introduction

The software architecture level of design enables us to cope with the increasing size and complexity of software systems during the early stage of their development [12, 13]. To achieve this, the design focus is turned from algorithmic and data structure related issues to the overall architecture of a system, where the architecture is meant to be a collection of computational components together with a description of their interactions. As software architecture emerges as a discipline within software engineering, it becomes increasingly important to support architectural devel-

opment with languages and tools. It is widely recognized that suitable architectural description languages (ADLs for short) should be devised to formalize software architectures instead of using informal box-and-line diagrams, and companion tools should be implemented to support the automatic analysis of architectural properties in order to allow the designer to make principled choices. Among the formal method based ADLs appeared in the literature, we mention those relying on process algebras [2, 10, 3], Z [1], and the CHAM [7].

The formal description of the software architecture of a system serves two purposes. First and foremost is making available a precise document describing the structure of the system to all the people involved in the design, implementation, and maintenance of the system itself. The second one is concerned with the possibility of analyzing the properties of the system at the architectural level, which allows for the early detection of errors thus resulting in time and money saving. This paper is about the latter purpose.

Nowadays systems are typically made out of numerous components whose behavior is individually well known. Think e.g. of embedded systems or web centered applications. Thus, the main problem faced by a software architect is that of understanding whether the components coordinate, i.e. fit together well. If the architecture of a system is given a formal description, then mathematical techniques can be used to automatically prove the well formedness of the system, i.e. the absence of architectural mismatches. In this paper we deal with three causes of architectural mismatches: incompatibility between two components due to a single interaction, incompatibility between two components due to the combination of several interactions, and lack of interoperability among a set of interacting components forming a cyclic topology.

The first cause is the simplest example of architectural mismatch: think e.g. of a component that wants to send data to another component via a certain connector, with

the connector never willing to accept data to be forwarded. This kind of architectural mismatch has been considered in the process algebraic framework of WRIGHT [2]. Using WRIGHT the architect is required to specify the ports of each component and the roles of each connector, where a port represents the behavior of a component w.r.t. an interaction with a connector, while the roles of a connector describe the expected interactions of each component attached to the connector. Then, given a connector and the set of components attached to it, all the pairs of corresponding ports and roles are verified to be compatible in order to ensure an overall deadlock free interaction of the set of components through that connector, under the assumption that each involved component has a single interaction with the connector.

In [8] it has been argued that mismatches should not be checked at the port/role level only, because the internal (as opposed to interactional) behavior of a component may correlate different interactions of the component itself in such a way that the system blocks. As an example, in [8] the CHAM has been employed to model a compressing proxy, whose gzip component interacts properly with the rest of the system as long as its input flow interaction and its output flow interaction are considered separately; however, since the gzip component can autonomously decide to start sending compressed data before having received all the data to be compressed, a deadlock is detected when considering the two flows together w.r.t. the rest of the system. This kind of architectural mismatch has been dealt with in [8, 6] by extracting from each component description its actual behavior and its assumptions about the behavior of the rest of the system, and by defining an algorithm that tries to match each assumed behavior with a sequence of actual behaviors. If the algorithm succeeds, the set of interacting components is deadlock free. In [3] the same kind of mismatch has been treated in a process algebraic framework by means of standard projections and observational equivalences. The idea is to verify that, for each pair composed of a component and a connector attached to each other, the behavior of the component projected on all the interactions with the connector is observationally equivalent to the behavior of the connector projected on all the interactions with the component. The first contribution of this paper is to strengthen the architectural compatibility check of [3] in order to achieve a deadlock freedom result, thus subsuming the compatibility check of [2] and equalling the effectiveness of the algorithm of [8, 6]. The importance of such a contribution is that different techniques for detecting similar architectural mismatches are brought into a uniform, process algebraic framework using standard machinery.

The second contribution of this paper is to show that a third cause of architectural mismatches can be revealed through standard observational equivalences in a process

algebraic framework. The point here is that mismatches should not be checked only between pairs of connected components, but also among sets of interacting components that form a cycle. The reason is that in such a scenario each pair composed of two interacting components could be well formed, while the global interaction might not be. In terms of the algorithm of [8, 6], the interactions within the sequence of actual behaviors matching an assumed behavior should be taken into account as well. As an example of this kind of mismatch, let us elaborate on the guest analogy of [8]. Here we have three components: the guest, the host, and the waiter. When arriving at the party, the guest expects to be welcome by the host and to be asked whether he/she wants an orange juice or a pineapple juice. The guest then expects that the host tells the waiter the desired drink and that the waiter brings that drink. Suppose that the binary interactions within the pairs of components are correct, but assume that the host has bad memory or is a malicious person and can thus tell the waiter to bring to the guest a drink different from the desired one. It is clear that the three components do not interoperate correctly when considered collectively. In this paper we define an architectural interoperability check in the framework of [3], which guarantees the absence of deadlock for sets of interacting components that form a cyclic topology.

This paper is organized as follows. In Sect. 2 and 3 we recall our formal framework based on process algebras, observational equivalences, and PADL [3]. In Sect. 4 we strengthen the architectural compatibility check of [3] so as to achieve a deadlock freedom result. In Sect. 5 we provide a novel architectural interoperability check in such a way that a deadlock freedom result is obtained. In Sect. 6 we consider the case of a cruise control system examined in [9] and we show that our architectural interoperability check is able to detect a well known mismatch. In Sect. 7 we report some concluding remarks.

2. Process Algebras

Process algebras [11, 5] are algebraic languages which support the compositional description of concurrent and distributed systems and the formal verification of their properties. The basic elements of any process algebra are its actions, which represent activities carried out by the systems being modeled, and its operators (among which a parallel composition operator), which are used to compose algebraic descriptions.

The set of process terms of PA is generated by the following syntax

$$E ::= \underline{0} \mid a.E \mid E/L \mid E[\varphi] \mid E + E \mid E \parallel_S E \mid A$$

where a belongs to a set Act of actions including a distinguished action τ for unobservable activities, $L, S \subseteq$

$a.E \xrightarrow{a} E$	
$\frac{E \xrightarrow{a} E'}{E/L \xrightarrow{a} E'/L} \text{ if } a \notin L$	$\frac{E \xrightarrow{a} E'}{E/L \xrightarrow{\tau} E'/L} \text{ if } a \in L$
$\frac{E \xrightarrow{a} E'}{E[\varphi] \xrightarrow{\varphi(a)} E'[\varphi]}$	
$\frac{E_1 \xrightarrow{a} E'}{E_1 + E_2 \xrightarrow{a} E'}$	$\frac{E_2 \xrightarrow{a} E'}{E_1 + E_2 \xrightarrow{a} E'}$
$\frac{E_1 \xrightarrow{a} E'_1}{E_1 \parallel_S E_2 \xrightarrow{a} E'_1 \parallel_S E_2} \text{ if } a \notin S$	$\frac{E_2 \xrightarrow{a} E'_2}{E_1 \parallel_S E_2 \xrightarrow{a} E_1 \parallel_S E'_2} \text{ if } a \notin S$
$\frac{E_1 \xrightarrow{a} E'_1 \quad E_2 \xrightarrow{a} E'_2}{E_1 \parallel_S E_2 \xrightarrow{a} E'_1 \parallel_S E'_2} \text{ if } a \in S$	
$\frac{E \xrightarrow{a} E'}{A \xrightarrow{a} E'} \text{ if } A \triangleq E$	

Table 1. Operational semantics for PA

$Act - \{\tau\}$, φ belongs to a set $ARFun$ of action relabeling functions preserving observability (i.e., $\varphi^{-1}(\tau) = \{\tau\}$), and A belongs to a set $Const$ of constants each possessing a (possibly recursive) defining equation of the form $A \triangleq E$.

In the syntax above, the null term “ $\underline{0}$ ” is the term that cannot execute any action. The prefix operator “ $a._$ ” denotes the sequential composition of an action and a term: term $a.E$ can execute action a and then behaves as term E . The hiding operator “ $_/L$ ” makes actions unobservable: term E/L behaves as term E except that each executed action a is turned into τ whenever $a \in L$. The relabeling operator “ $_{[\varphi]}$ ” changes actions: term $E[\varphi]$ behaves as term E except that each executed action a is turned into $\varphi(a)$. The alternative composition operator “ $_ + _$ ” expresses a nondeterministic choice between two terms: term $E_1 + E_2$ behaves as either term E_1 or term E_2 depending on whether an action of E_1 or an action of E_2 is executed. The parallel composition operator “ $_ \parallel_S _$ ” expresses the concurrent execution of two terms according to the following synchronization discipline: two actions can synchronize if and only if they are equal and observable. Term $E_1 \parallel_S E_2$ asynchronously executes actions of E_1 or E_2 not belonging to S and syn-

chronously executes actions of E_1 and E_2 belonging to S if the requirement above is met. The prefix operator and the alternative composition operator are called dynamic operators, whereas the hiding operator, the relabeling operator, and the parallel composition operator are called static operators. We say that a term is sequential if it is composed of dynamic operators only.

The semantics for PA is defined in the standard operational style by means of a set of axioms and inference rules, which formalize the meaning of each operator and produce as semantic model a state transition graph, where states are in correspondence with process terms and transitions are labeled with actions. The semantics for PA is shown in Table 1. As an example, the first rule for the hiding operator means that, if the current state is E/L and E is capable of performing an action $a \notin L$ thus evolving into E' , then E/L is capable of performing the same action a thus evolving into E'/L . The other rules must be read in the same way. In particular, note that in the third rule for the parallel composition operator the action resulting from the synchronization of two a actions is still a . This means that multiway synchronizations are allowed.

Due to their algebraic nature, process description languages like PA naturally lend themselves to the definition of equivalences and preorders. In particular, in the concurrency theory literature several notions of equivalence can be found, which relate terms according to a certain interpretation of their behavior. Here we recall the weak bisimulation equivalence, which captures the ability of two terms to simulate each other behavior up to τ actions. Before introducing it, we generalize the transition relation \longrightarrow labeled with actions to the transition relation \Longrightarrow labeled with sequences of actions: $\xrightarrow{a_1 \dots a_n} \equiv \xrightarrow{a_1} \dots \xrightarrow{a_n}$, with ε denoting the empty string and $\xrightarrow{\varepsilon}$ being the identity relation over process terms. Moreover, if σ is a sequence over Act , let $\hat{\sigma}$ denote the sequence over $Act - \{\tau\}$ obtained from σ by removing all the occurring τ actions. Finally, we use $\xRightarrow{\sigma} \equiv \xrightarrow{\tau^m} \xrightarrow{\sigma} \xrightarrow{\tau^n}$, with $m, n \in \mathbb{N}$, to indicate the execution of an action sequence σ possibly preceded and followed by the execution of arbitrarily many invisible actions.

A binary relation \mathcal{B} over process terms is a weak bisimulation [11] iff, whenever $(E_1, E_2) \in \mathcal{B}$, then for all $a \in Act$:

- whenever $E_1 \xrightarrow{a} E'_1$, then $E_2 \xRightarrow{\hat{a}} E'_2$ and $(E'_1, E'_2) \in \mathcal{B}$ for some E'_2 ;
- whenever $E_2 \xrightarrow{a} E'_2$, then $E_1 \xRightarrow{\hat{a}} E'_1$ and $(E'_1, E'_2) \in \mathcal{B}$ for some E'_1 .

Note that $\hat{a} = a$ for $a \neq \tau$, $\hat{a} = \varepsilon$ for $a = \tau$. The union of all the weak bisimulations, denoted by \approx_B , is called the weak bisimulation equivalence. \approx_B enjoys several algebraic properties w.r.t. the dynamic operators:

$$\begin{aligned}
(E_1 + E_2) + E_3 &\approx_B E_1 + (E_2 + E_3) \\
E_1 + E_2 &\approx_B E_2 + E_1 \\
E + \underline{0} &\approx_B E \\
E + E &\approx_B E \\
\tau.E &\approx_B E \\
a.\tau.E &\approx_B a.E \\
E + \tau.E &\approx_B \tau.E \\
a.(E_1 + \tau.E_2) + a.E_2 &\approx_B a.(E_1 + \tau.E_2)
\end{aligned}$$

The last four properties show the ability of \approx_B to abstract from τ actions and make \approx_B suited to reason on projections (obtained through the hiding operator) of the behavior of system components. Furthermore, \approx_B is a congruence w.r.t. the static operators. This means that the observable behavior of a term does not change if we replace a subterm with a weakly bisimulation equivalent term in the context of a static operator. Formally, if $E_1 \approx_B E_2$ then

$$\begin{aligned}
E_1/L &\approx_B E_2/L \\
E_1[\varphi] &\approx_B E_2[\varphi] \\
E_1 \parallel_S E &\approx_B E_2 \parallel_S E
\end{aligned}$$

This will be exploited in the proof of Thm. 4.3.

3. PADL

A description in PADL, the ADL based on PA introduced in [3], is an architectural type (AT). Each AT is defined as a function of its architectural element types (AETs), its topology, and its architectural interactions (AIs). An AET is in turn defined as a function of its behavior, specified either as a family of PA sequential terms or through an invocation of a previously specified AT, and its interactions, specified as a set of PA actions. The architectural topology consists of a fixed set of architectural element instances (AEIs) related by a fixed set of attachments among their interactions. Finally, AIs are interactions of AEIs that act as interactions for the whole AT, thus supporting hierarchical architectural modeling.

An architectural description in PADL is as shown in Table 2. We now illustrate the meaning of each part of the architectural description above by means of an example concerning a pipe-filter system. It is composed of three identical filters and a pipe. Each filter acts as a service center with a two position buffer. For each item processed by the upstream filter, the pipe forwards it to one of the two downstream filters according to the availability of free positions in their buffers. If both have free positions, the choice is resolved nondeterministically.

The first part of a PADL description defines the name of the AT and its formal parameters, i.e. AETs and AIs:

archi_type *PipeFilter*(*FilterT*, *PipeT*; *PipeFilterI*)

The specification above indicates that the AT named *PipeFilter* relies on two AETs, *FilterT* and *PipeT*, and on a set *PipeFilterI* of AIs.

The second part of the description defines the AETs that have been specified to be formal parameters of the AT in the first part. Each AET is defined through a family of PA sequential terms (or an invocation of a previously specified AT) expressing its behavior and a set of PA actions expressing its interactions:

elem_type *FilterT*(*Filter*; *FilterI*)

behavior *Filter* \triangleq *accept_item.Filter'*
Filter' \triangleq *accept_item.Filter''* +
serve_item.Filter
Filter'' \triangleq *serve_item.Filter'*

interactions *FilterI* = *accept_item*, *serve_item*

elem_type *PipeT*(*Pipe*, *PipeI*)

behavior *Pipe* \triangleq *accept_item*.
(*forward_item*₁.*Pipe* +
*forward_item*₂.*Pipe*)
interactions *PipeI* = *accept_item*,
*forward_item*₁,
*forward_item*₂

The specification above indicates that AET *FilterT* has the

archi_type	⟨name and formal parameters⟩
archi_elem_types	⟨architectural element types: behaviors and interactions⟩
archi_topology	⟨architectural topology: architectural element instances and attachments⟩
archi_interactions	⟨architectural interactions⟩
end	

Table 2. PADL textual notation

behavior described by term *Filter*, which represents a service center of capacity two, and the interactions described by action set *FilterI*. Initially, the filter can only accept an item coming from the outside, thus evolving into *Filter'*. If an item is in the buffer (*Filter'*), then either another item from the outside is accepted or the item in the buffer is served. If two items are in the buffer (*Filter''*), no more items can be accepted; the filter can only serve one of the waiting items. Similarly, AET *PipeT* has the behavior described by term *Pipe* and the interactions described by action set *PipeI*. The pipe repeatedly accepts an item and forwards it along one of two different routes.

The third part of the description defines the topology of the AT in terms of instances of the previously introduced AETs and related attachments among their interactions:

archi_topology $F_0, F_1, F_2 : FilterT$
 $P : PipeT$
 $F_0.serve_item \text{ to } P.accept_item$
 $F_1.accept_item \text{ to } P.forward_item_1$
 $F_2.accept_item \text{ to } P.forward_item_2$

The attachments above specify that an instance P of *PipeT* is connected to an upstream instance F_0 of *FilterT* and two downstream instances F_1, F_2 of *FilterT*.

Finally, the fourth part defines the AIs. They are interactions of the previously defined AEIs that are used when defining the behavior of an AET of another AT by invoking the AT being specified, i.e. when plugging the AT at hand in the context of a larger system:

archi_interactions $PipeFilterI = F_0.accept_item,$
 $F_1.serve_item,$
 $F_2.serve_item$

This part supports hierarchical modeling. For instance, the pipe-filter architecture may represent the structure of the server in a client-server architecture.

PipeFilter can be pictorially represented through the flow graph [11] of Fig. 1, where boxes denote AEIs, black circles/white squares denote interactions/AIs, and edges between ports denote attachments.

The semantics of a PADL specification is given by translation into PA as follows.




Figure 1. Flow graph of *PipeFilter*

Definition 3.1 Given a PADL specification, let C be an AET with behavior E and interactions I . The semantics of C and its instances is defined to be its behavior projected onto its interactions:

$$\llbracket C \rrbracket = E / (Act - \{\tau\} - I) \quad \blacksquare$$

In our example we have $\llbracket FilterT \rrbracket = \llbracket F_0 \rrbracket = \llbracket F_1 \rrbracket = \llbracket F_2 \rrbracket = Filter / (Act - \{\tau\} - FilterI)$ and $\llbracket PipeT \rrbracket = \llbracket P \rrbracket = Pipe / (Act - \{\tau\} - PipeI)$.

The semantics of an AT is then obtained by composing in parallel the semantics of its AEIs according to the specified attachments. Recalled that the parallel composition operator is left associative, in our example we have:

$$\begin{aligned} \llbracket PipeFilter \rrbracket = & \llbracket F_0 \rrbracket [serve_item \mapsto a] \parallel \emptyset \\ & \llbracket F_1 \rrbracket [accept_item \mapsto a_1] \parallel \emptyset \\ & \llbracket F_2 \rrbracket [accept_item \mapsto a_2] \parallel \{a, a_1, a_2\} \\ & \llbracket P \rrbracket [accept_item \mapsto a] \\ & \quad [forward_item_1 \mapsto a_1, \\ & \quad \quad forward_item_2 \mapsto a_2] \end{aligned}$$

The use of the relabeling operator is necessary to make AEIs interact. As an example, F_0 and P must interact via *serve_item* and *accept_item*, which are different from each other. Since the parallel composition operator allows actions to synchronize only if they are equal, in $\llbracket PipeFilter \rrbracket$ above each *serve_item* executed by $\llbracket F_0 \rrbracket$ and

each *accept_item* executed by $\llbracket P \rrbracket$ is relabeled to the same action *a* (occurring neither in $\llbracket F_0 \rrbracket$ nor in $\llbracket P \rrbracket$ to avoid interferences), then a synchronization on *a* is forced between the relabeled versions of $\llbracket F_0 \rrbracket$ and $\llbracket P \rrbracket$ by means of operator $\llbracket \{a, a_1, a_2\} \cdot$.

To define the semantics of an arbitrary AET, we preliminarily introduce some notation that will be used also in the rest of the paper. Let $\{C_1, \dots, C_n\}$ be a subset of its AEIs, with interactions I_{C_1}, \dots, I_{C_n} respectively; we use i, j, k to range over $\{1, \dots, n\}$. We say that $C_i.a_1$ is connected to $C_j.a_2$ iff either there is an attachment between them, or there exists an interaction a_3 of C_k such that $C_i.a_1$ is connected to $C_k.a_3$ and there is an attachment between $C_k.a_3$ and $C_j.a_2$. Moreover, we say that a subset of attachments is connected iff the involved interactions are pairwise connected via attachments of the subset only and the subset is maximal. Denoted by $I_{C_i; C_1, \dots, C_n} \subseteq I_{C_i}$ the subset of interactions of C_i attached to C_1, \dots, C_n , let $S(C_1, \dots, C_n)$ be a set of as many fresh¹ actions as there are connected subsets of attachments among the considered AEIs of the PADL specification, let $\varphi_{C_i; C_1, \dots, C_n} : I_{C_i; C_1, \dots, C_n} \rightarrow S(C_1, \dots, C_n)$ be action relabeling functions such that $\varphi_{C_i; C_1, \dots, C_n}(a_1) = \varphi_{C_j; C_1, \dots, C_n}(a_2)$ iff $C_i.a_1$ is connected to $C_j.a_2$, and let $S(C_i; C_1, \dots, C_n) = \varphi_{C_i; C_1, \dots, C_n}(I_{C_i; C_1, \dots, C_n})$ and $S(C_i, C_j; C_1, \dots, C_n) = S(C_i; C_1, \dots, C_n) \cap S(C_j; C_1, \dots, C_n)$.

Definition 3.2 Given the PADL specification above, we define the interacting semantics $\llbracket C_i \rrbracket_{C_1, \dots, C_n}$ of C_i restricted to C_1, \dots, C_n by

$$\llbracket C_i \rrbracket_{C_1, \dots, C_n} / (Act - \{\tau\} - I_{C_i; C_1, \dots, C_n}) [\varphi_{C_i; C_1, \dots, C_n}]$$

The semantics $I(C_1, \dots, C_n)$ of the considered group of AEIs is given by

$$\llbracket C_1 \rrbracket_{C_1, \dots, C_n} \parallel_{S(C_1, C_2; C_1, \dots, C_n)} \dots \parallel_{\bigcup_{i=1}^{n-1} S(C_i, C_n; C_1, \dots, C_n)} \llbracket C_n \rrbracket_{C_1, \dots, C_n} \quad \blacksquare$$

The PA term above is the semantics of a whole PADL specification whenever the considered group comprises all of its AEIs, in which case the hiding occurring in the definition of the restricted interacting semantics is redundant.

We conclude by observing that PADL is closely related to WRIGHT [2]. In particular, we note that the behavior of AETs is specified through process terms, and that the semantics is given by translation into a process term by essentially exploiting the parallel composition operator. However, there are several differences between PADL and WRIGHT. First, WRIGHT distinguishes between components and connectors, while in PADL there are just architectural elements, each of which can be interpreted as being a component or a connector depending on the particular system. This avoids redundancy in the specifications that results in the presence of connectors whose behavior is

trivial. Second, in WRIGHT the description of each component/connector is accompanied by its ports/roles, whereas in PADL the interactions are simply expressed as actions (similarly to DARWIN [10]). This allows ports and roles to be retrieved via projections (see the definition of restricted interacting semantics) and avoids redundancy in the specifications. Third, PADL allows the hierarchical modeling of families of architectures and provides an efficient technique to check an architecture instance for conformity to an architectural type [3]. Fourth, PADL is equipped with a graphical notation that is standard in the concurrency theory literature.

4. Compatibility Check

In order to make sure that every pair of AEIs attached to each other interact in a proper way, in [3] PADL has been equipped with an architectural compatibility check. It consists of verifying that, for each pair of AEIs attached to each other, their restricted interacting semantics are weakly bisimulation equivalent. According to the notation of Def. 3.2, we say that C_i is compatible with C_j iff $\llbracket C_i \rrbracket_{C_i, C_j} \approx_B \llbracket C_j \rrbracket_{C_i, C_j}$. In our example this condition is met by all the pairs of interacting AEIs:

$$\begin{aligned} \llbracket F_0 \rrbracket / (Act - \{\text{serve_item}, \tau\}) [\text{serve_item} \mapsto a] &\approx_B \\ \llbracket P \rrbracket / (Act - \{\text{accept_item}, \tau\}) [\text{accept_item} \mapsto a] & \\ \llbracket F_1 \rrbracket / (Act - \{\text{accept_item}, \tau\}) [\text{accept_item} \mapsto a_1] &\approx_B \\ \llbracket P \rrbracket / (Act - \{\text{forward_item}_1, \tau\}) [\text{forward_item}_1 \mapsto a_1] & \\ \llbracket F_2 \rrbracket / (Act - \{\text{accept_item}, \tau\}) [\text{accept_item} \mapsto a_2] &\approx_B \\ \llbracket P \rrbracket / (Act - \{\text{forward_item}_2, \tau\}) [\text{forward_item}_2 \mapsto a_2] & \end{aligned}$$

In general such a compatibility check does not guarantee that the interaction of the AEIs attached to a given AEI is deadlock free. In this section we consider only ATs whose reduced flow graph (in which all the attachments between two AEIs are collapsed into a single one) is acyclic; cycles will be dealt with in the next section. The absence of cycles implies that all the AEIs interacting through a given AEI cannot interact in other ways.

Definition 4.1 Given an acyclic PADL specification, let C_1, \dots, C_n be the AEIs attached to AEI K . The interaction $I(K; C_1, \dots, C_n)$ of C_1, \dots, C_n through K is defined by

$$\begin{aligned} \llbracket K \rrbracket_{K, C_1, \dots, C_n} \parallel_{S(K; K, C_1)} \\ \llbracket C_1 \rrbracket_{K, C_1, \dots, C_n} \parallel_{S(K; K, C_2)} \dots \\ \dots \parallel_{S(K; K, C_n)} \llbracket C_n \rrbracket_{K, C_1, \dots, C_n} \end{aligned}$$

$I(K; C_1, \dots, C_n)$ is deadlock free iff for each of its states s there exist a visible action $a \in Act - \{\tau\}$ and a state s' such that $s \xrightarrow{a} s'$. \blacksquare

Deadlock freedom holds for the interactions of systems like *PipeFilter*, but if we take e.g. $\llbracket C \rrbracket \triangleq \tau.C_1 + \tau.C_2, C_1 \triangleq$

¹I.e., not occurring in the semantics of the considered AEIs.

$a.C_1, C_2 \triangleq b.C_2$ and $\llbracket K \rrbracket \triangleq \tau.K_1 + \tau.K_2, K_1 \triangleq a.K_1, K_2 \triangleq b.K_2$ with equal actions attached to each other and identically relabeled, then $\llbracket C \rrbracket_{K,C} \approx_B \llbracket K \rrbracket_{K,C}$ but $I(K; C) \equiv \llbracket K \rrbracket_{K,C} \parallel_{\{a,b\}} \llbracket C \rrbracket_{K,C}$ is not deadlock free even if $\llbracket C \rrbracket_{K,C}$ and $\llbracket K \rrbracket_{K,C}$ are such.

A sufficient condition for deadlock freedom can be retrieved by slightly modifying the architectural compatibility check of [3] as follows. Let us define the interacting semantics of AEI C_i w.r.t. AEI C_j by $\llbracket C_i \rrbracket'_{C_i, C_j} = \llbracket C_i \rrbracket[\varphi_{C_i, C_j}]$.

Definition 4.2 Given an acyclic PADL specification, let C_i and C_j be two AEIs. C_j is said to be compatible with C_i iff

$$\llbracket C_i \rrbracket'_{C_i, C_j} \parallel_{S(C_i, C_i, C_j)} \llbracket C_j \rrbracket_{C_i, C_j} \approx_B \llbracket C_i \rrbracket'_{C_i, C_j} \quad \blacksquare$$

Theorem 4.3 Given an acyclic PADL specification, let C_1, \dots, C_n be the AEIs attached to AEI K . If $\llbracket K \rrbracket$ is deadlock free and C_i is compatible with K for all $i = 1, \dots, n$, then $I(K; C_1, \dots, C_n)$ is deadlock free.

Proof The result follows by proving that $I(K; C_1, \dots, C_n) \approx_B \llbracket K \rrbracket_{K, C_1, \dots, C_n}$ as $\llbracket K \rrbracket_{K, C_1, \dots, C_n}$ is deadlock free, because it differs from $\llbracket K \rrbracket$ only for the names of its interactions, and \approx_B preserves deadlock freedom. This can be proved by induction on n by observing that $\llbracket K \rrbracket'_{K, C_i} \parallel_{S(K; K, C_i)} \llbracket C_i \rrbracket_{K, C_i} \approx_B \llbracket K \rrbracket'_{K, C_i} \implies \llbracket K \rrbracket_{K, C_1, \dots, C_n} \parallel_{S(K; K, C_i)} \llbracket C_i \rrbracket_{K, C_i} \approx_B \llbracket K \rrbracket_{K, C_1, \dots, C_n}$ for all $i = 1, \dots, n$, because \approx_B is a congruence w.r.t. relabeling, and by exploiting the fact that \approx_B is a congruence w.r.t. parallel composition. \blacksquare

It is worth noting that verifying architectural compatibility over each pair C_i, K is more convenient than globally verifying that $I(K; C_1, \dots, C_n)$ is deadlock free. The first reason is that the local check requires the construction of n state spaces resulting from the parallel composition of two terms only (see Def. 4.2), whereas the global check requires the construction of one state space resulting from the parallel composition of n terms (see Def. 4.1). Thus, the complexity of the local check is linear in the number of AEIs, while the complexity of the global check is exponential in the number of AEIs. The second reason is that, in case of mismatch, the local check provides the architect with diagnostic information to single out the AEIs responsible for the mismatch, while the global check does not.

We conclude by observing that this new architectural compatibility check not only strengthens that of [3], but also subsumes the compatibility check of [2] and equals the effectiveness of the algorithm of [8, 6], thus bringing into a uniform, process algebraic framework different techniques for detecting similar architectural mismatches using standard machinery.

5. Interoperability Check

Checking architectural compatibility is unfortunately not enough to guarantee that a system is deadlock free, i.e. there may be further causes of architectural mismatches. Let us consider e.g. the flow graph of Fig. 2, which represents the guest analogy mentioned in Sect. 1. Although the AEIs are deadlock free, the whole system is not. Here the problem is that *Guest* and *Waiter* do not communicate directly, so e.g. if *Guest* asks for an orange juice and *Host* tells *Waiter* to bring a pineapple juice instead, *Guest* and *Waiter* do not synchronize, thus causing the system to block. If we construct the state space, we discover that there are two deadlock states arising from the situation just mentioned and its symmetrical.

If we look carefully at the flow graph of Fig. 2, we observe that the three AEIs form a loop. The architectural mismatch arises because *Host* can act as a malicious agent between *Guest* and *Waiter*. Should the behavior of *Host* and *Waiter* be concentrated in a single AEI interacting with *Guest*, the mismatch would be detected through the architectural compatibility check.

In order to catch this kind of mismatches, in this section we define a suitable architectural interoperability check.

Definition 5.1 Given a PADL specification, let C_1, \dots, C_n be AEIs forming a loop. C_1, \dots, C_n are said to interoperate iff for all $i = 1, \dots, n$

$$I(C_1, \dots, C_n) / (Act - \{\tau\} - S(C_i; C_1, \dots, C_n)) \approx_B \llbracket C_i \rrbracket_{C_1, \dots, C_n} \quad \blacksquare$$

Theorem 5.2 Given a PADL specification, let C_1, \dots, C_n be AEIs forming a loop. If there exists $i \in \{1, \dots, n\}$ such that $\llbracket C_i \rrbracket_{K_1, \dots, K_m}$ is deadlock free and C_1, \dots, C_n interoperate, then $I(C_1, \dots, C_n)$ is deadlock free.

Proof A straightforward consequence of the fact that \approx_B preserves deadlock. \blacksquare

Unlike the architectural compatibility check, the architectural interoperability check guarantees the absence of deadlock with a cost that is exponential in the number of AEIs in the loop. One cannot even think of reducing the cost by modifying the architectural interoperability check by e.g. comparing each AEI of the loop with the overall loop interaction with that AEI excluded, otherwise the problem about deadlock freedom discussed in the first part of Sect. 4 shows up. It is however worth noting that resorting to the architectural interoperability check is more convenient than globally checking $I(C_1, \dots, C_n)$ for deadlock. The first reason is that, while performing the former check, the state space of $I(C_1, \dots, C_n)$ can be reduced via \approx_B thanks to the use of the hiding operator in Def. 5.1. The second reason is that the former check provides diagnostic information in case of mismatch. For instance, in the scenario of

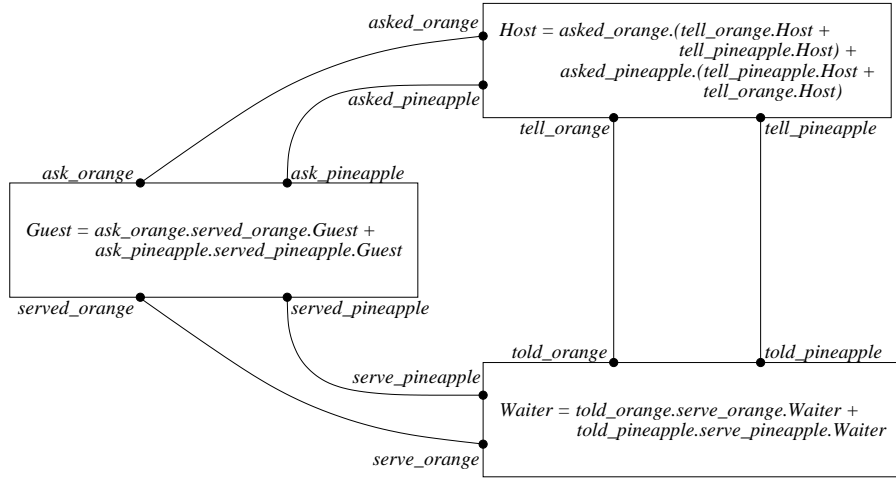


Figure 2. Flow graph representation of the guest analogy

Fig. 2, we find that *Guest* fails the architectural interoperability check. By exploiting the modal logic based diagnostic information provided by tools like [4] in case of equivalence checking failure, we know that if *Guest* performs e.g. *ask_orange* then it necessarily performs *served_orange* afterwards, while this is not the case for the whole loop. Since by direct inspection we see that *Guest* is correct, we know that the cause of the mismatch is in the rest of the loop. In general, when a loop is formed by many AEIs, if one of such AEIs fails the architectural interoperability check, but by direct inspection we see that it and its adjacent AEIs are correct, the loop can be shrunk by replacing those AEIs with their parallel composition and the procedure is repeated until the cause of the architectural mismatch is identified.

6. Cruise Control System

In this section we assess the validity of our novel architectural interoperability check to discover a known architectural mismatch in the cruise control system examined in [9].

An automobile cruise control system is controlled by three buttons: on, off, and resume. When the engine is running and on is pressed, the cruise control system records the current speed and maintains the car at this speed. When the accelerator, brake, or off is pressed, the cruise control system disengages but retains the speed setting. If resume is pressed, the system accelerates or decelerates the car back to the previously recorded speed.

The PADL description of the system is depicted in Fig. 3. The behavior of each AEI is shown in Table 3. The system has 84 states and 230 transitions.

It can be globally verified by tools like [4] that this system is deadlock free, the reason being that speed measurement related interactions take place endlessly as long as

the engine is running. Indeed, if we hide those interactions, the presence of deadlock shows up. The existence of a problem is revealed by the fact that, when considering the loop of Fig. 3, *Sensor* fails the architectural interoperability check. By exploiting the modal logic based diagnostic information stemming from the failure of equivalence checking, we discover that the interaction of the loop in Fig. 3 deadlocks after performing *turn_engine_on*, *press_on*, *turn_engine_off*, and *turn_engine_on* again, while *Sensor* does not, and that all we have to do in order to overcome this problem is to insert action *trigger_disable_speed_control* after *turned_engine_off* in *Cruising*. It is worth observing that we have been able to automatically and systematically discover such a subtle mismatch thanks to our architectural interoperability check, whereas in [9] it is mentioned that the same error went originally undetected for some years and was discovered for the first time while animating the specification.

7. Conclusion

In this paper we have strengthened the architectural compatibility check of [3], thus subsuming the compatibility check of [2] and equalling the effectiveness of the algorithm of [8, 6] in a uniform framework thanks to standard observational equivalences, and we have introduced a novel architectural interoperability check, whose adequacy has been assessed on a cruise control system. Both checks, which rely on the properties of the observational equivalences defined for process algebras, come equipped with a deadlock freedom result, which makes them more convenient – for efficiency and/or diagnosis related reasons – than globally verifying deadlock freedom. A lesson learnt is that process algebras are well suited at the architectural level of design

<i>Sensor</i>	\triangleq	<i>turn_engine_on.Sensor'</i>
<i>Sensor'</i>	\triangleq	<i>press_accelerator.Sensor' +</i> <i>press_brake.Sensor' +</i> <i>press_on.Sensor' +</i> <i>press_off.Sensor' +</i> <i>press_resume.Sensor' +</i> <i>turn_engine_off.Sensor</i>
<i>CruiseController</i>	\triangleq	<i>Inactive</i>
<i>Inactive</i>	\triangleq	<i>turned_engine_on.trigger_clear_speed.Active</i>
<i>Active</i>	\triangleq	<i>pressed_accelerator.Active +</i> <i>pressed_brake.Active +</i> <i>pressed_on.trigger_enable_speed_control.trigger_record_speed.Cruising +</i> <i>pressed_off.Active +</i> <i>pressed_resume.Active +</i> <i>turned_engine_off.Inactive</i>
<i>Cruising</i>	\triangleq	<i>pressed_accelerator.trigger_disable_speed_control.StandBy +</i> <i>pressed_brake.trigger_disable_speed_control.StandBy +</i> <i>pressed_on.trigger_enable_speed_control.trigger_record_speed.Cruising +</i> <i>pressed_off.trigger_disable_speed_control.StandBy +</i> <i>pressed_resume.Cruising +</i> <i>turned_engine_off.Inactive</i>
<i>StandBy</i>	\triangleq	<i>pressed_accelerator.trigger_disable_speed_control.StandBy +</i> <i>pressed_brake.trigger_disable_speed_control.StandBy +</i> <i>pressed_on.trigger_enable_speed_control.trigger_record_speed.Cruising +</i> <i>pressed_off.trigger_disable_speed_control.StandBy +</i> <i>pressed_resume.trigger_enable_speed_control.Cruising +</i> <i>turned_engine_off.Inactive</i>
<i>SpeedDetector</i>	\triangleq	<i>turned_engine_on.WheelRevCounter</i>
<i>WheelRevCounter</i>	\triangleq	<i>measure_speed.signal_speed.WheelRevCounter +</i> <i>turned_engine_off.SpeedDetector</i>
<i>SpeedController</i>	\triangleq	<i>Disabled</i>
<i>Disabled</i>	\triangleq	<i>signalled_speed.Disabled +</i> <i>triggered_clear_speed.clear_speed.Disabled +</i> <i>triggered_enable_speed_control.enable_speed_control.Enabled +</i> <i>triggered_disable_speed_control.disable_speed_control.Disabled</i>
<i>Enabled</i>	\triangleq	<i>signalled_speed.maintain_speed.adjust_throttle.Enabled +</i> <i>triggered_enable_speed_control.enable_speed_control.Enabled +</i> <i>triggered_record_speed.record_speed.Enabled +</i> <i>triggered_disable_speed_control.disable_speed_control.Disabled</i>
<i>Throttle</i>	\triangleq	<i>adjusted_throttle.Throttle</i>

Table 3. Component behaviors for the cruise control system

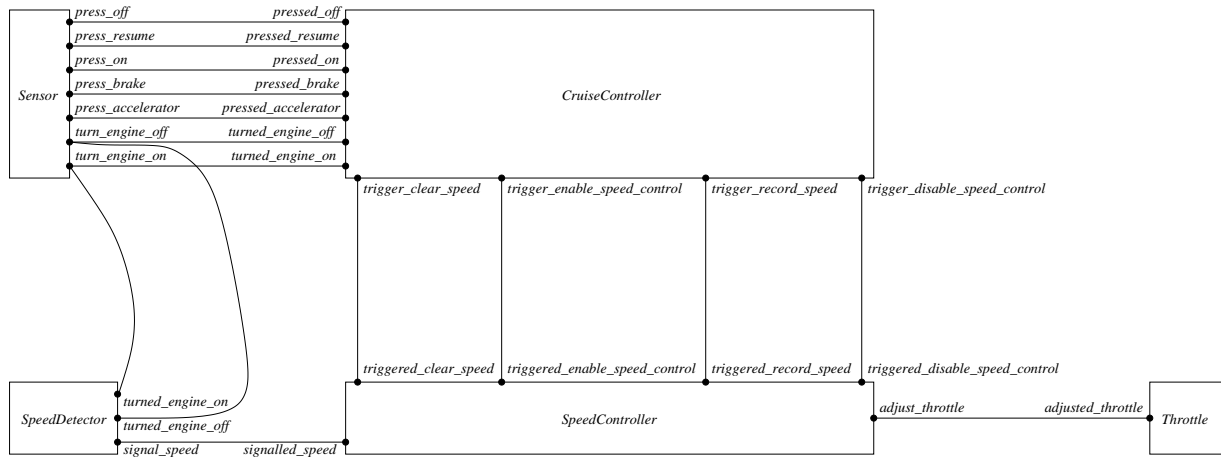


Figure 3. Flow graph representation of the cruise control system

not only because of their compositional modeling facilities, but also for their analysis related machinery that enables us to detect architectural mismatches in a rather efficient way and to get some diagnostic information useful for identifying the mismatch sources.

As far as future work is concerned, first we would like to understand to which extent the loop shrinking procedure sketched at the end of Sect. 5 can be automated (to reduce the architect intervention). Second, we would like to weaken the architectural compatibility and interoperability checks using other observational equivalences and investigate to which extent the deadlock freedom results are still valid. Different observational equivalences may be adequate to detect architectural mismatches that can be characterized by properties other than deadlock. Third, we would like to understand whether our deadlock freedom results, which are local to AEIs and loop interactions, can be generalized to the whole AT. In other words, we would like to see whether there are further causes of architectural mismatches, and whether they can be managed by local checks.

Acknowledgements

This research has been funded by Progetto MURST Cofinanziato SALADIN.

References

- [1] G.D. Abowd, R. Allen, D. Garlan, "Formalizing Style to Understand Descriptions of Software Architecture", in ACM Trans. on Software Engineering and Methodology 4:319-364, 1995
- [2] R. Allen, D. Garlan, "A Formal Basis for Architectural Connection", in ACM Trans. on Software Engineering and Methodology 6:213-249, 1997
- [3] M. Bernardo, P. Ciancarini, L. Donatiello, "On the Formalization of Architectural Types with Process Algebras", in Proc. of the 8th ACM Int. Symp. on the Foundations of Software Engineering (FSE-8), ACM Press, pp. 140-148, San Diego (CA), 2000
- [4] W.R. Cleaveland, J. Parrow, B. Steffen, "The Concurrency Workbench: A Semantics-Based Tool for the Verification of Concurrent Systems", in ACM Trans. on Programming Languages and Systems 15:36-72, 1993
- [5] C.A.R. Hoare, "Communicating Sequential Processes", Prentice Hall, 1985
- [6] P. Inverardi, S. Uchitel, "Proving Deadlock Freedom in Component-Based Programming", in Proc. of the 4th Int. Conf. on Fundamental Approaches to Software Engineering (FASE '01), LNCS 2029:60-75, Genova (Italy), 2001
- [7] P. Inverardi, A.L. Wolf, "Formal Specification and Analysis of Software Architectures Using the Chemical Abstract Machine Model", in IEEE Trans. on Software Engineering 21:373-386, 1995
- [8] P. Inverardi, A.L. Wolf, D. Yankelevich, "Static Checking of System Behaviors Using Derived Component Assumptions", in ACM Trans. on Software Engineering and Methodology 9:239-272, 2000
- [9] J. Kramer, J. Magee, "Exposing the Skeleton in the Coordination Closet", in Proc. of the 2nd Int. Conf. on Coordination Languages and Models (COORDINATION '97), LNCS 1282:18-31, Berlin (Germany), 1997
- [10] J. Magee, N. Dulay, S. Eisenbach, J. Kramer, "Specifying Distributed Software Architectures", in Proc. of the 5th European Software Engineering Conf. (ESEC '95), LNCS 989:137-153, Barcelona (Spain), 1995
- [11] R. Milner, "Communication and Concurrency", Prentice Hall, 1989
- [12] D.E. Perry, A.L. Wolf, "Foundations for the Study of Software Architecture", in ACM SIGSOFT Software Engineering Notes 17:40-52, 1992
- [13] M. Shaw, D. Garlan, "Software Architecture: Perspectives on an Emerging Discipline", Prentice Hall, 1996