

Unblockable Compositions of Software Components

Ruzhen Dong
UNU-IIST, Macau
ruzhen@iist.unu.edu

Jiří Srba
Dept. of Computer Science
Aalborg University, Denmark
srba@cs.aau.dk

Johannes Faber
UNU-IIST, Macau
jfaber@iist.unu.edu

Naijun Zhan
State Key Lab. of Comp. Sci.
Institute of Software, CAS
znpj@ios.ac.cn

Zhiming Liu
UNU-IIST, Macau
lzm@iist.unu.edu

Jiaqi Zhu
State Key Lab. of Comp. Sci.
Institute of Software, CAS
zhujq@ios.ac.cn

ABSTRACT

We present a new automata-based interface model describing the interaction behavior of software components. Contrary to earlier component- or interface-based approaches, the composition of components in our model guarantees that no behavior of one component can be blocked by the other component, independent of the actual implementation of the component as long as the interface description is respected. To this end, we develop an algorithm to compute the unblockable interaction behavior, called the *interface model of a component*, from its execution model. Based on this model, we introduce composition and coordination operators for the components, and we prove important compositionality results, showing the conditions under which composition of interface models preserves unblockable sequences of provided services.

1. INTRODUCTION

In component-based software engineering, large software systems are decomposed into individual software components with clearly articulated interfaces in order to facilitate a sound development process across different teams of developers. An interface theory should then define the basic principles for composing several software components based on their interfaces while the concrete implementation of the components is invisible to its environment. This means that components can be treated as back-boxes and the theory allows for independent implementability.

Since the 90s, component-based approaches have drawn a lot of attention in software engineering [Szy97, AG97, Arb04, LMS09] and there has been a considerable research activity studying how interfaces of components can be formally described in order to automatically decide whether two components can be composed together in a well-formed way [LT89, dAH01b, DAH01a, LNW07a, DLL⁺10]. However, none of these approaches can guarantee that the composition of several components does not lead to malfunction-

ing components exposed to potentially blocked behaviours.

We examine to what extent an interface model can be realized in a way that ensures nonblocking executions for all possible compositions of components respecting the given interfaces, and present an automata-based interface model that possesses this desirable property. In our model, a component comprises of a *provided interface* and a *required interface*. The former describes which executions of services the component offers to its environment, while the latter specifies what services the component needs to call in order to provide the services from its provided interface. The *execution model* of a component describes the interaction behavior between the component and its environment.

The *interface model* we propose in this paper supports a black-box composition in the sense that components can be composed without an a priori knowledge of their concrete implementation as long as the specified sequences of provided services and the corresponding sequences of required services are respected by the implementation. This in particular means that any sequence of services specified in the provided interface cannot be blocked when being composed with any other component. We prove that this is the case if and only if the interface model is *input-deterministic*, meaning that after calling any fixed sequence of provided services, there is a unique selection of currently available provided services, and this selection is independent of the component's internal state and, in particular, also independent of the service sequences the component has called in order to execute its provided services.

In order to generate an interface model for any given execution model, we describe a new algorithm that computes all unblockable sequences of services provided by a component from its execution model. For the application of our model in the context of a component-based design approach, we also introduce composition and coordination operators for manipulating the components. The *composition* of components is used to synchronize services provided by one and required by another component. A specific form of composition is the *plugging* of components, where one component only provides services to the other component without requiring any services from it. On the other hand, the services of components can be *coordinated* by triggering services internally via a coordination process. We prove that unblockable sequences of services provided by the plugging of execution models are the same as those of the plugging of their interface models, and show why similar results do not hold for composition and coordination.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

Comparison with I/O and Interface Automata.

There are two broadly-known approaches to interface theories, the I/O Automata [LT89, LT87] and the Interface Automata [dAH01b, DAH01a, DAH05]. Both are centered around an interface-based design and have some similarities to our approach, however, they have been developed with different aims and are based on different assumptions. In fact, we shall argue that our approach is positioned between these existing approaches, trying to overcome their limitations.

Input/Output (I/O) Automata [LT89, LT87] were defined by Lynch and Tuttle to model concurrent and distributed discrete event systems. The actions are separated into input, output, and internal actions. In this model, input-enabledness is required in the sense that all the input actions must be enabled at any state. Synchronization is realized by broadcasting on the outputs, i.e., an output synchronizes with all corresponding input actions of the other automata. The compatibility of two components is checked in a *pessimistic* way, which requires that there is no deadlock for *all* possible environments. On the contrary, our interface model does not enforce that all inputs are always enabled which can be sometimes unrealistic to achieve, but we assume input-determinism in order to guarantee that all sequences of provided services accepted by the interface model can be executed without even being blocked.

To overcome the restrictions of I/O Automata, which are forced to react in a pessimistic way to every possible input, de Alfaro and Henzinger introduced their Interface Automata [dAH01b, DAH01a, DAH05] that are based on the *optimistic* assumption that a component is normally used in a specific environment and thus does not need to react to every input. In this model, methods of a component that can be called are modeled as input actions; external methods that the component calls are output actions; internal actions are used to model internal method invocations. Interface automata guarantee the specified output executions on the assumption that the environment only invokes methods provided as the inputs. The optimistic compatibility of two interface automata implies that two interface automata are compatible if there is at least one environment avoiding deadlock states in composed product of the two components. Such a condition for compositionally may be too relaxed for certain applications.

In contrast to the existing approaches, the model we present in this paper combines both the optimistic and pessimistic views for component-based systems: the interface model of a component directly specifies which sequences of provided services guarantee deadlock-free executions. Thus, it is on the one hand not as pessimistic as I/O automata, because all provided services do not have to be available at every state, and on the other hand, it is not as optimistic as interface automata, because we restrict the possible service-invocation environments such that, under these restricted environments, we have a guarantee of non-blocking behavior in any possible execution.

In the rest of the paper, we shall analyze the properties of our new interface model and demonstrate in which situations it is superior to the previous models and in which cases using the existing techniques is beneficial.

Summary of contributions.

The contributions of this paper are (1) a new interface

model ensuring unblockable compositions of software components; (2) an algorithm to generate the interface model of a component based on its execution model; (3) definitions of basic operations used in component-based design; (4) a formal analysis showing that unblockable behavior is preserved for the full plugging of components but not in the case of general composition and coordination.

Outline of the paper.

The rest of the paper is organized as follows. In Sect. 2, we introduce component automata and an algorithmic way to generate their interface models. In Sect. 3, we present the composition operators to compose components as well as processes and prove important properties for these operators. In Sect. 4, we show how components can be coordinated by active processes, and in Sect. 5 we conclude the paper and discuss the future work.

2. COMPONENT AUTOMATA

In this section, we will first introduce some notions that will be used throughout the paper and then motivate component automata and component interface automata. At last, we will give an algorithm transforming component automata to component interface automata and prove its correctness.

2.1 Preliminary Definitions

Let $A, B \subseteq \mathcal{L}^*$ be two languages. The set concatenation $A \circ B$ is defined as $\{w_A w_B \mid w_A \in A, w_B \in B\}$. The first and the second projection on a pair of elements $\ell = (x, y)$ are denoted by π_1 and π_2 , and defined by $\pi_1(\ell) = x$ and $\pi_2(\ell) = y$, and it is naturally extended to sequences of pairs of elements. For a sequence $\rho = \langle x_0, x_1, \dots, x_k \rangle$ over an alphabet \mathcal{L} and a set $X \subseteq \mathcal{L}$, a projection on the elements from X is denoted by $\rho|X$ and defined as $\langle x_{i_1}, x_{i_2}, \dots, x_{i_n} \rangle$ where $0 \leq i_1 < i_2 < \dots < i_n \leq k$ are all the indices of elements such that $x_{i_j} \in X$ for all j , $1 \leq j \leq n$. The empty sequence is denoted as ϵ . The concatenation of sequences tr_1 and tr_2 is denoted by $tr_1 \circ tr_2$, and also extended to sets of sequences such that $T_1 \circ T_2$ contains all concatenations $tr_1 \circ tr_2$ for $tr_1 \in T_1$ and $tr_2 \in T_2$.

2.2 Execution Model of a Component

An execution of a component can be modeled as an alternating sequence of provided service invocations (initiated by the environment) and a set of causality sequences of required service invocations (initiated by the component). In our automata-based model, the invocation of a provided service is modeled as an atomic *provided event* and the invocation of the required services as a sequence of atomic *required events*. As there may be many choices of required services in order to provide a service, we allow sets of sequences of required events with the intuition that an implementation of the component may choose which of the sequences it actually implements in order to provide a service. The formal definition is as follows.

DEFINITION 1. A tuple $C = (S, s_0, P, R, \delta)$ is called a component automaton where

- ◇ S is a finite set of states,
- ◇ $s_0 \in S$ is the initial state,
- ◇ P and R are disjoint and finite sets of provided and required events, respectively,

$\diamond \delta \subseteq S \times \Sigma(P, R) \times S$ is the transition relation, where the set of labels is defined as $\Sigma(P, R) = P \times (2^{R^*} \setminus \emptyset)$.

Whenever there is $(s, \ell, s') \in \delta$ with $\ell = (a, T)$, we simply write $s \xrightarrow{a/T} s'$ and call it a transition step. A component automaton is called *closed* if $T = \{\epsilon\}$ for all transitions $s \xrightarrow{a/T} s'$. Otherwise, it is called *open*. In component design, one is finally interested in closed components, i.e., components which are not depending on further services from the environment.

An alternating sequence of states and labels of the form

$$e = \langle s_0, \ell_0, s_1, \ell_1, \dots, s_k, \ell_k, s_{k+1} \rangle$$

is called an *execution* of the component automaton C if $s_i \xrightarrow{\ell_i} s_{i+1}$ for all i , $0 \leq i \leq k$, where s_0 is the initial state. A sequence of labels $tr = \langle \ell_0, \ell_1, \dots, \ell_k \rangle$ is called a *trace* of C if there is an execution e of C such that $tr = e|_{\Sigma(P, R)}$. In other words, a trace contains only the labels of an execution and abstracts away from the states. The set of all traces of C is denoted as $\mathcal{T}(C)$, and we write $s_0 \xRightarrow{tr} s$ if there is an execution from s_0 to s under the trace tr . A sequence of provided events $pt \in P^*$ is called a *provided trace* if there is a trace tr such that $\pi_1(tr) = pt$. The set of all provided traces of C is denoted as $\mathcal{T}_p(C)$. For a given state s of a component $C = (S, s_0, P, R, \delta)$ the set of provided traces starting from this state s is defined by $\mathcal{T}_p((S, s, P, R, \delta))$ and denoted as $\mathcal{T}_p(s)$.

Given a provided trace $pt \in \mathcal{T}_p(C)$, the set *caused*(pt) of caused traces for pt is defined as

$$\begin{aligned} \text{caused}(pt) &= \{T_0 \circ \dots \circ T_k \mid tr \in \mathcal{T}(C), \\ &\quad \pi_1(tr) = pt, \pi_2(tr) = \langle T_0, \dots, T_k \rangle\} . \end{aligned}$$

EXAMPLE 1. As a demonstrating example, we consider a simple internet-connection component presented in Fig. 1. It provides the services *login*, *wifi*, *print*, *read*, and *disc* available to the environment. The services model the logging into the system, request for wifi connection, invocation of printing a document, an email service, and disconnecting from the internet, respectively. The component calls the services *unu1*, *unu2*, *cserv*, *getm*. The first three of them model the searching for a wifi router nearby, connecting to the *unu1* or *unu2* wireless network, and connecting to an application server, respectively. The *getm* is a service that fetches an email, connects to a printer, sends a document for a print and starts the printing job. The *print* service is only available for the wifi network *unu1* and *read* can be accessed at both networks.

In the component model of Fig. 1 we can perform e.g.

$$e = \langle 0, (\text{login}/\{\epsilon\}), 1, (\text{wifi}/\{\text{unu1}\}), 2, (\text{print}/\{\epsilon\}), 2 \rangle .$$

Now $pt = \langle \text{login}, \text{wifi}, \text{print} \rangle$ is a provided trace of the execution e and the set of caused traces of pt is $\text{caused}(pt) = \{\text{unu1}\}$. This example will be used throughout this paper to show the features of our model.

2.3 Unblockable Equivalence

The environment interacts with the component in a way that it chooses a sequence of provided events and the component (non-deterministically) determines the sequence of required events that it will request from other components.

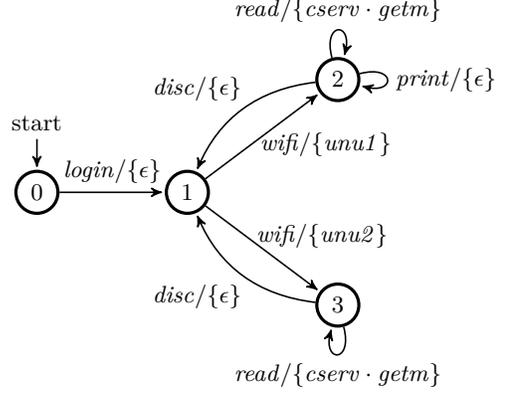


Figure 1: Execution model of internet connection component C_{ic}

As the interface of a component should be defined independent of the actual internal choices done by the component, we have to make sure that if the interface model provides a sequence of events in one execution scenario, it is also capable to do so in any other execution scenario determined by the component. This guarantees that a possible component composition will be deadlock free, regardless of the actual implementation of the components.

In the example, we may observe that including the provided event *print* to the interface of the internet connection component from Fig. 1 is not safe, as the sequence of requested services $\langle \text{login}, \text{wifi}, \text{print} \rangle$ can be executed if the connection is made to the network *unu1*, but it will be impossible if network *unu2* was chosen. This will cause a blocking in the component composition and as the environment is not in control of the selection of the network, exposing *print* to the component interface model is not safe; the service should not be advertised. For this reason, we will require that our interface models of components have the property of *input-determinism*, meaning that after any sequence of provided service invocations, the component should be at states with identical provided services. This will guarantee that the provided traces are not blocked during any run-time execution.

Given a component automaton $C = (S, s_0, P, R, \delta)$, let us define for any $s \in S$ the set of enabled provided events as

$$\text{enabled}(s) = \{a \in P \mid s \xrightarrow{a/T} s'\} .$$

Then, we define the concept of input-determinism.

DEFINITION 2 (INPUT-DETERMINISM). A component automaton is input-deterministic iff

$$s_0 \xRightarrow{tr_1} s_1 \text{ and } s_0 \xRightarrow{tr_2} s_2$$

with $\pi_1(tr_1) = \pi_1(tr_2)$ implies

$$\text{enabled}(s_1) = \text{enabled}(s_2) .$$

Now we can also define the notion of an unblockable trace. An unblockable provided trace is never blocked by the component in the sense that the component can always provide the trace independent of how the internal choices of the causality traces are resolved in an implementation of the component.

DEFINITION 3 (UNBLOCKABLE TRACE). We call a provided trace $pt = \langle a_0, \dots, a_k \rangle$ of a component automaton $C = (S, s_0, P, R, \delta)$ unblockable iff whenever $s_0 \xrightarrow{tr} s$ such that $\pi_1(tr) = \langle a_0, \dots, a_i \rangle$ for some i , $0 \leq i \leq k-1$, then $a_{i+1} \in \text{enabled}(s)$. A trace in C is unblockable iff its provided trace is unblockable. The set of all unblockable traces of C is denoted by $\mathcal{T}_u(C)$.

The next theorem shows that the property of input-determinism guarantees that all traces are unblockable and vice versa.

THEOREM 1. A component automaton $C = (S, s_0, P, R, \delta)$ is input-deterministic iff all of its traces are unblockable.

PROOF. First, we prove the direction from left to right. From the input-determinism of C follows that for each provided trace $pt = \langle a_0, \dots, a_k \rangle$ and for each state s with $s_0 \xrightarrow{tr} s$ and $\pi_1(tr) = \langle a_0, \dots, a_i \rangle$ for $0 \leq i \leq k-1$, the set $\text{enabled}(s)$ is the same. Since pt is a provided trace, so $a_{i+1} \in \text{enabled}(s)$. This shows that all provided traces are unblockable, and so all traces are unblockable too.

Second, we prove the direction from right to left by contraposition. We assume that C is not input-deterministic, so there exist traces tr_1 and tr_2 with $\pi_1(tr_1) = \pi_1(tr_2)$ and $s_0 \xrightarrow{tr_1} s_1$, $s_0 \xrightarrow{tr_2} s_2$ such that $\text{enabled}(s_1) \neq \text{enabled}(s_2)$. Then we can w.l.o.g. assume that there is a provided event a such that $a \in \text{enabled}(s_1)$ and $a \notin \text{enabled}(s_2)$. Now $\pi_1(tr_1) \circ \langle a \rangle$ is a provided trace of C that is blockable. \square

For the purpose of our interface theory, we will consider two component automata equivalent whenever their sets of unblockable traces are the same.

DEFINITION 4 (UNBLOCKABLE EQUIVALENCE). Let C_1 and C_2 be two component automata. They are equivalent with respect to unblockable traces, denoted by $C_1 \equiv C_2$, if $\mathcal{T}_u(C_1) = \mathcal{T}_u(C_2)$. Similarly, we write $C_1 \equiv_{pt} C_2$ if the sets of unblockable provided traces of C_1 and C_2 are equal.

2.4 Interface Model of a Component

In the example shown in Fig. 1, the provided trace $\langle \text{login}, \text{wifi}, \text{print} \rangle$ may be blocked, when the component resolves the non-determinism by choosing to call `unu2`. This reveals that the component contains a potential deadlock when composed with another component using the print service. Thus, we propose an input-deterministic model to support deadlock-free black-box composition as the interface model.

DEFINITION 5. An interface automaton is an input deterministic component automaton.

We now present an algorithm that, for a given component automaton C , constructs the largest interface automaton $\mathcal{I}(C)$ that is equivalent with respect to its unblockable traces with the component automaton. The automaton $\mathcal{I}(C)$ can be understood as the interface behavior of the component that can be advertised to other components and guarantees a deadlock-free composition as long as the interaction described by the interface automaton is respected.

A general construction of an interface automaton $\mathcal{I}(C)$ for a given component automaton C is given in Algorithm 1. The states of the interface automaton are pairs (Q, r) consisting of a subset of states Q of the original automaton and

Algorithm 1: Construction of Interface Automaton $\mathcal{I}(C)$

Input: $C = (S, s_0, P, R, \delta)$
Output: $\mathcal{I}(C) = (S_I, (\{s_0\}, s_0), P, R, \delta_I)$, where $S_I \subseteq 2^S \times S$

- 1: **Initialization:** $S_I := \{(\{s_0\}, s_0)\}$; $\delta_I := \emptyset$;
 $todo := \{(\{s_0\}, s_0)\}$; $done := \emptyset$
- 2: **while** $todo \neq \emptyset$ **do**
- 3: **choose** $(Q, r) \in todo$;
 $todo := todo \setminus \{(Q, r)\}$; $done := done \cup \{(Q, r)\}$
- 4: **for each** $a \in \bigcap_{s \in Q} \text{enabled}(s)$ **do**
- 5: $Q' := \bigcup_{s \in Q} \{s' \mid (s \xrightarrow{a/T} s') \in \delta\}$
- 6: **for each** $(r \xrightarrow{a/T} r') \in \delta$ **do**
- 7: **if** $(Q', r') \notin (todo \cup done)$ **then**
- 8: $todo := todo \cup \{(Q', r')\}$
- 9: $S_I := S_I \cup \{(Q', r')\}$
- 10: **end if**
- 11: $\delta_I := \delta_I \cup \{(Q, r) \xrightarrow{a/T} (Q', r')\}$
- 12: **end for**
- 13: **end for**
- 14: **end while**

a single state r . In the first element of the pair, the power-set construction (similar to the construction of a deterministic automaton from a non-deterministic one) is adopted to identify the set of all states Q that can be reached by a provided trace in C , which is used to identify all non-blocking provided traces. In the second element of the pair, we compute the intersection of the languages of provided traces of C and the constructed power-set automaton. By this, the original executions of the C are simulated for the non-blockable provided traces.

EXAMPLE 2. We illustrate our algorithm on an example shown in Fig. 2. All states of $\mathcal{I}(C)$ consist of pairs. In the first element of the pair, we have the set of states reachable by the executions in C , and a state has a transition with a provided event e if and only if all states in the first element of the pair enable such a transition. For instance, the state $(\{1, 2\}, 1)$ has an outgoing transition labeled with a_2 , because in C both states 1 and 2 enable a_2 . The second element of the state pair, namely 1, has influence on the caused trace: the outgoing transition of $(\{1, 2\}, 1)$ is labeled with T_4 , because the a_2 transition of state 1 in C is also labeled with T_4 . On the other hand, there is no outgoing transition labeled with a_3 for $(\{1, 2\}, 1)$ as a_3 is enabled only from the state 2 but not from the state 1.

We now argue about the correctness of the algorithm. In what follows, we fix a component automaton $C = (S, s_0, P, R, \delta)$ and the interface automaton $\mathcal{I}(C) = (S_I, (\{s_0\}, s_0), P, R, \delta_I)$ produced by Algorithm 1. We will not explicitly mention whether the transitions are performed in C or $\mathcal{I}(C)$ as this is clear from the structure of the states. We begin with three lemmas showing the relation between the states (Q, r) and r (Lemma 1), showing that unblockable traces in C are preserved in $\mathcal{I}(C)$ (Lemma 2), and showing that traces in $\mathcal{I}(C)$ correspond to unblockable traces in C (Lemma 3).

LEMMA 1. Let $(Q, r) \in S_I$ be a reachable state in the interface automaton $\mathcal{I}(C)$.

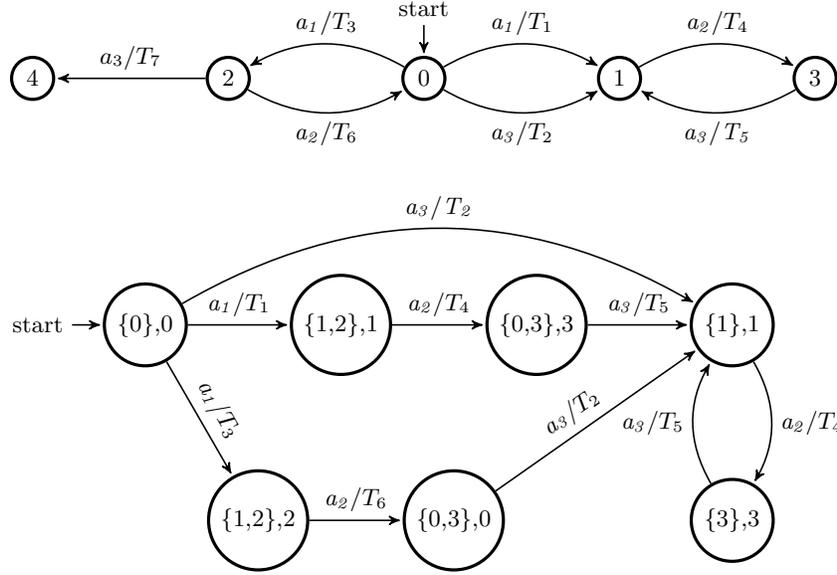


Figure 2: Component automaton C top, interface automaton $\mathcal{I}(C)$ bottom

1. If $r \xrightarrow{a/T} r'$ such that $a \in \bigcap_{s \in Q} \text{enabled}(s)$ then

$$(Q, r) \xrightarrow{a/T} (Q', r')$$

where $Q' = \{s' \mid s \xrightarrow{a/T'} s', s \in Q\}$.

2. If $(Q, r) \xrightarrow{a/T} (Q', r')$ then

$$Q' = \{s' \mid s \xrightarrow{a/T'} s', s \in Q, T' \in 2^{R^*}\} \text{ and } r \xrightarrow{a/T} r'.$$

PROOF. For the first claim assume that (Q, r) is a reachable state in $\mathcal{I}(C)$, implying that (Q, r) was during some iteration of the while-loop selected at line 3 of the algorithm. From the assumption $a \in \bigcap_{s \in Q} \text{enabled}(s)$ we get that the body of the for-loop at line 4–13 will be executed for such an a . The set Q' is then constructed as required at line 5 and the transition $(Q, r) \xrightarrow{a/T} (Q', r')$ is added at line 11 thanks to the assumption that $r \xrightarrow{a/T} r'$.

For the second claim we again observe that

$$(Q, r) \xrightarrow{a/T} (Q', r') \quad (1)$$

implies that $Q' = \{s' \mid s \xrightarrow{a/T'} s', s \in Q, T' \in 2^{R^*}\}$ as such Q' is constructed at line 5, and (1) is added at line 11. Clearly, the transition can be added only if $r \xrightarrow{a/T} r'$ because this is the condition of the for-loop at line 6–12. \square

LEMMA 2. Let $pt \in \mathcal{T}_p(C)$ be an unblockable provided trace in C . If $s_0 \xrightarrow{tr} r$ where $\pi_1(tr) = pt$, then

$$(\{s_0\}, s_0) \xrightarrow{tr} (Q, r)$$

such that $Q = \{s' \mid s_0 \xrightarrow{tr'} s', \pi_1(tr') = pt\}$.

PROOF. By induction on the length of pt : The base case where $|pt| = 0$ follows directly from the definitions and line 1 of the algorithm. Let $|pt| > 0$. Then $pt = pt' \cdot a$ for some provided event a . Let $tr = tr' \cdot a/T$ be a trace with

$\pi_1(tr) = pt$ such that $s_0 \xrightarrow{tr} r$ and $r \xrightarrow{a/T} r'$. By induction hypothesis, $(\{s_0\}, s_0) \xrightarrow{tr'} (Q, r)$ where

$$Q = \{s' \mid s_0 \xrightarrow{tr''} s', \pi_1(tr'') = pt'\}.$$

As the provided trace $pt' \cdot a$ is by assumption unblockable, we have $a \in \bigcap_{s \in Q} \text{enabled}(s)$. By the fact that $r \xrightarrow{a/T} r'$ and Lemma 1 we get $(Q, r) \xrightarrow{a/T} (Q', r')$ where

$$Q' = \{s' \mid s \xrightarrow{a/T'} s', s \in Q\},$$

so in other words $Q' = \{s \mid s_0 \xrightarrow{tr''} s, \pi_1(tr'') = pt\}$ and $(\{s_0\}, s_0) \xrightarrow{tr} (Q', r')$. \square

LEMMA 3. If $(\{s_0\}, s_0) \xrightarrow{tr} (Q, r)$ then

$$Q = \{s \mid s_0 \xrightarrow{tr'} s, \pi_1(tr') = \pi_1(tr)\}$$

and $s_0 \xrightarrow{tr} r$, where tr is unblockable in C .

PROOF. By induction on the length of tr . The base case where $|tr| = 0$ is trivial. Let $|tr| > 0$. Then $tr = tr' \cdot a/T$ such that $(\{s_0\}, s_0) \xrightarrow{tr'} (Q, r)$ and $(Q, r) \xrightarrow{a/T} (Q', r')$. By induction hypothesis,

$$Q = \{s \mid s_0 \xrightarrow{tr''} s, \pi_1(tr'') = \pi_1(tr')\}$$

and $s_0 \xrightarrow{tr'} r$ such that tr' is unblockable in C . By the fact that $(Q, r) \xrightarrow{a/T} (Q', r')$ and Lemma 1, we have

$$Q' = \{s' \mid s \xrightarrow{a/T'} s', s \in Q\} \text{ and } r \xrightarrow{a/T} r'.$$

By combining these observations we get $s_0 \xrightarrow{tr} r'$ and also $Q' = \{s \mid s_0 \xrightarrow{tr''} s, \pi_1(tr'') = \pi_1(tr' \cdot a/T)\}$. Notice that by the construction of the interface automaton $\mathcal{I}(C)$ we have $a \in \bigcap_{s \in Q} \text{enabled}(s)$, so $tr' \cdot a/T$ is unblockable since tr' is unblockable by the induction hypothesis. \square

Now we can present an important theorem stating that the constructed interface automaton preserves the set of unblockable traces of the input component automaton.

THEOREM 2. *Algorithm 1 terminates for every given component automaton C and produces a unique interface automaton $\mathcal{I}(C)$ such that $C \equiv \mathcal{I}(C)$.*

PROOF. The termination follows from the fact that in the interface automaton there are only finitely many possible states of the form (Q, r) . Every iteration of the while-loop at line 2–14 removes one element from the set *todo* and places it to the set *done* (line 3) and any (Q, r) can be inserted to the set *todo* at most once due to the condition at line 7. Also, even though the choice of (Q, r) at line 3 is non-deterministic, the algorithm returns the same output $\mathcal{I}(C)$ as each reachable state in the interface automaton has to be processed, independent of the order which is chosen.

We prove that $\mathcal{I}(C)$ is input-deterministic. By Lemma 3 we get that for any two states in the interface automaton $\mathcal{I}(C)$ that are reachable with the same provided trace, the first element of the state pair is equal: (Q, r_1) and (Q, r_2) . Thus, we only need to show that we have $enabled(Q, r_1) = enabled(Q, r_2)$. However, this follows from $enabled(Q, r) = \bigcap_{s \in Q} enabled(s)$.

Thus, $\mathcal{I}(C)$ is an interface automaton, and hence all of its traces are unblockable. Lemma 2 shows that $\mathcal{T}_u(C) \subseteq \mathcal{T}_u(\mathcal{I}(C))$ and Lemma 3 shows that $\mathcal{T}_u(\mathcal{I}(C)) \subseteq \mathcal{T}_u(C)$. Hence $C \equiv \mathcal{I}(C)$. \square

COROLLARY 1. *Whenever $C_1 \equiv C_2$ for two component automata C_1 and C_2 then also $\mathcal{I}(C_1) \equiv \mathcal{I}(C_2)$.*

PROOF. Follows directly from Theorem 2 and from the fact that the unblockable equivalence \equiv is transitive. \square

3. COMPOSITION OPERATORS

“Components are for composition.” [Szy97]. A component interacts with other components by providing services or requiring services. A closed component provides services without the need to require services from other components. We consider closed components as stable service providers. Open components will provide services under assumptions that the required services are guaranteed. Composition allows for components interacting with each other to build up new components. It also enables reusability and decomposition of components. In this section, we will introduce some basic composition operators including the product and the plugging of components.

3.1 Product of Component Automata

The product of two components assembles the components together and synchronizes them when a sequence of required service calls can be provided by the other component. We call two components composable if their provided services are disjoint, because invocations to joint services will trigger both of the components, which will cause non-deterministic executions.

Let us now motivate the product composition of component automata. The definition is similar to the standard product for finite automata, except that we synchronize transitions if a caused service of one component is provided by a corresponding transition of the other component, e.g., for transitions

$$r_1 \xrightarrow{a/\{b_1 \cdot b_2\}} r'_1 \text{ and } r_2 \xrightarrow{b_1/\{c_1 \cdot c_2\}} r'_2$$

in C_1 and C_2 , the synchronized transition will be

$$(r_1, r_2) \xrightarrow{a/\{c_1 \cdot c_2 \cdot b_2\}} (r'_1, r'_2).$$

In definition of the product, there is an additional condition for synchronization. Assume that there also exists

$$r_1 \xrightarrow{a/\{d_1\}} r''_1$$

where d_1 is a provided event of C_2 but $d_1 \notin enabled(r_2)$. With this, the component C_1 can trigger the service d_1 that is declared as provided in C_2 , but component C_2 cannot provide it in r_2 . This causes a potential deadlock in the component C_1 , because whether $b_1 \cdot b_2$ or d_1 is called is determined by the component internally, in other words, it is invisible to the outside. The solution to avoid the potential internal deadlock is to make a unavailable at state (r_1, r_2) . Informally, should the product construction guarantee a at state (r_1, r_2) , it must require that all caused traces of a at state r_1 restricted to the set of provided services of C_2 are also provided by C_2 at state r_2 .

We use the notation $caused(s, a) = \{\alpha \mid s \xrightarrow{a/T} s', \alpha \in T\}$ to denote the set of all caused traces for the provided event a at state s ; by $events(T)$ we denote the set of required events that occur in the set T of required event sequences.

DEFINITION 6 (PRODUCT). *Two component automata $C_1 = (S_1, s_0^1, P_1, R_1, \delta_1)$ and $C_2 = (S_2, s_0^2, P_2, R_2, \delta_2)$ are composable if $P_1 \cap P_2 = \emptyset$. The product $C_1 \otimes C_2$ is then defined as a component automaton (S, s_0, P, R, δ) where*

- $\diamond S = S_1 \times S_2$,
- $\diamond s_0 = (s_0^1, s_0^2)$,
- $\diamond P = (P_1 \setminus R_2) \cup (P_2 \setminus R_1)$,
- $\diamond R = (R_1 \setminus P_2) \cup (R_2 \setminus P_1)$, and
- $\diamond \delta$ is the smallest relation constructed as follows.

Let $r_1 \xrightarrow{a/T} r'_1 \in \delta_1$ with $a \in P$ and $r_2 \in S_2$. If there is a sequence $\alpha \in caused(r_1, a)$ such that $\alpha \setminus P_2 \notin \mathcal{T}_P(r_2)$ —that is, a service from the caused traces in r_1 which is provided by the second component is blocked by the possible executions starting in r_2 —then there is no corresponding a transition in δ . Otherwise we distinguish between a synchronization step and a non-synchronized step.

1. *Non-synchronized transition step:*
if $events(T) \cap P_2 = \emptyset$, then

$$(r_1, r_2) \xrightarrow{a/T} (r'_1, r_2) \in \delta.$$

2. *Synchronized transition step:*
if $events(T) \cap P_2 \neq \emptyset$, then

$$(r_1, r_2) \xrightarrow{a/T'} (r'_1, r'_2) \in \delta$$

whenever $r_2 \xrightarrow{\langle (a_0, T_0), \dots, (a_k, T_k) \rangle} r'_2$ with

$$\langle a_0, \dots, a_k \rangle = \beta \setminus P_2, \beta \in T$$

and

$$T' = \{\beta[\beta_0/a_0, \beta_1/a_1, \dots, \beta_k/a_k] \mid \beta_i \in T_i, 0 \leq i \leq k\}.$$

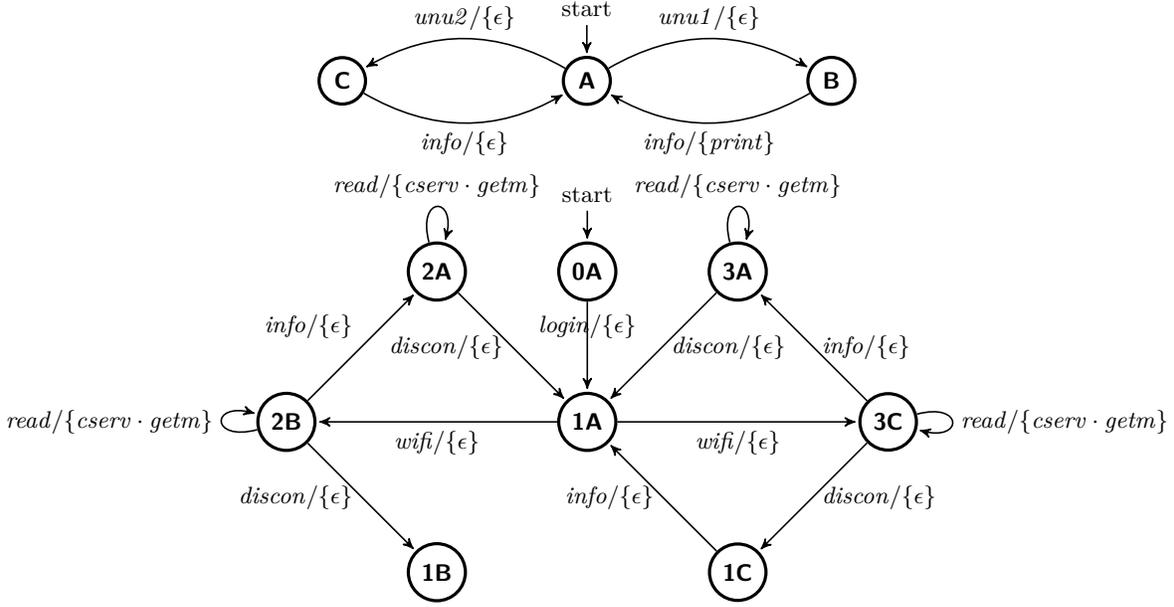


Figure 3: Component automaton C_{ip} top and product $C_{ic} \otimes C_{ip}$ bottom

These presented rules are based on transitions from δ_1 ; the rules for transitions from δ_2 are symmetric.

We will now use an example to illustrate the construction.

EXAMPLE 3. In the context of the internet component of Example 1, we consider a new component that provides the services *unu1*, *unu2*, *info* and requires the service *print*. The service *info* returns information on the status of this component. The execution model of this internet-provider component is pictured in the upper part of Fig. 3. The component at the bottom shows the product of the internet-connection component C_{ic} from Fig. 1 and the internet-provider component. The common services that are used for synchronization are *unu1*, *unu2*, *print*.

Note that the trace $\langle \text{login}, \text{wifi}, \text{info} \rangle$ is an unblockable trace of the product (there are two possible executions, one ending in 2A and one ending in 3A). Thus, this trace is also in the interface $\mathcal{I}(C_{ic} \otimes C_{ip})$. However, if we compute the interfaces of the product constituents, we will find that $\mathcal{I}(C_{ip})$ (equal to C_{ip} of Fig. 3) requires the *print* service. On the other hand, we saw in Sect. 2.3 that the interface of the internet-connection component C_{ic} does not contain the *print* service, even though the component provides this service. Thus, the trace $\langle \text{login}, \text{wifi}, \text{info} \rangle$ in the composition of the interfaces $\mathcal{I}(C_{ic}) \otimes \mathcal{I}(C_{ip})$ (that is the same as the product $C_{ic} \otimes C_{ip}$ without the state 2A) is blockable. Consequently, the product of two components can generally resolve potentially blocking executions, or in other words, the product of two components can have more unblockable executions than the unblockable traces produced by its constituents. Since the interface only contains the unblockable executions of the model and it is in general not possible to derive the interface of a product from the interfaces of the single components.

For this reason, we introduce a special case of product, the plugging of components, where the second component does

not required any events provided by the first one. This plugging product has the desired congruence property that the interface model of the composed components can be derived from the interfaces of its constituents.

DEFINITION 7 (PLUGGING). A component $C_1 = (S_1, s_0^1, P_1, R_1, \delta_1)$ is pluggable (can be plugged) by a component automaton $C_2 = (S_2, s_0^2, P_2, R_2, \delta_2)$ if C_1 and C_2 are composable and $P_1 \cap R_2 = \emptyset$. The plugging of C_1 by C_2 , denoted by $C_1 \ll C_2$, is then given as the product $C_1 \otimes C_2$.

We shall now discuss the properties of the plugging operator. The lemma below shows that a provided trace of a component automaton C_1 is preserved in the product with a component automaton C_2 if its causality traces projected to the services provided by C_2 are unblockable in C_2 .

In the rest of this section we fix two components $C_1 = (S_1, s_0^1, P_1, R_1, \delta_1)$ and $C_2 = (S_2, s_0^2, P_2, R_2, \delta_2)$ such that C_1 is pluggable by C_2 . Let $C_1 \ll C_2 = (S, s_0, P, R, \delta)$.

LEMMA 4. Let C_1 be pluggable by C_2 . If $s_0^1 \xrightarrow{tr_1} r_1$ in C_1 and $s_0^2 \xrightarrow{tr_2} r_2$ in C_2 with $\pi_1(tr_2) \in Q$ where $Q = \{\alpha \mid P_2 \mid \alpha \in \pi_2(tr_1)\}$ and all traces in Q are unblockable in C_2 , then $(s_0^1, s_0^2) \xrightarrow{tr} (r_1, r_2)$ for some tr with $\pi_1(tr) = \pi_1(tr_1)$.

PROOF. By induction on the length of tr_1 . The base case where $|tr_1| = 0$ is trivial. Let $|tr_1| > 0$. Then $tr_1 = tr'_1 \cdot a/T$ such that $s_0^1 \xrightarrow{tr'_1} r'_1$ and $r'_1 \xrightarrow{a/T} r_1$. Now $s_0^2 \xrightarrow{tr'_2 \cdot tr_2} r_2$ such that $s_0^2 \xrightarrow{tr'_2} r'_2$ and $r'_2 \xrightarrow{tr_2} r_2$ with that $\pi_1(tr'_2) \in Q'$ and $\pi_1(tr'_2 \cdot tr_2) \in Q$ where $Q' = \{\alpha \mid P_2 \mid \alpha \in \pi_2(tr'_1)\}$ and $Q = \{\alpha \mid P_2 \mid \alpha \in \pi_2(tr_1)\}$ are unblockable in C_2 . By induction hypothesis, there exists $(s_0^1, s_0^2) \xrightarrow{tr'} (r'_1, r'_2)$ where $\pi_1(tr') = \pi_1(tr'_1)$. By the transition rule of the product, we have $(r'_1, r'_2) \xrightarrow{a/T'} r_1, r_2$, because that Q and Q' are unblockable implies *caused*(r'_1, a) can be accepted by C_2 without being blocked. So $(s_0^1, s_0^2) \xrightarrow{tr} r_1, r_2$ where $\pi_1(tr) = \pi_1(tr') \cdot a = \pi_1(tr_1)$. \square

For the preservation of unblockable traces in plugging, it is obvious that the general plugging will not work. If there are some services provided by C_2 available also to the environment, the behavior of C_2 is uncertain and the unblockable traces cannot be preserved from C_1 to $C_1 \ll C_2$. Therefore, we consider a special case of plugging, *complete plugging*, requiring that all services of C_2 are provided only to C_1 and not available to the environment, formally $P_2 \subseteq R_1$. Such complete plugging is often the case in the component-based design when the functionality of a component is to solely provide services to a single component invoking the services.

THEOREM 3. *Let C_1 be pluggable by C_2 such that $P_2 \subseteq R_1$. Then a provided trace pt is an unblockable in $C_1 \ll C_2$ if and only if pt is unblockable in C_1 and all traces in $\{\alpha | P_2 \mid \alpha \in \text{caused}(pt)\}$ are unblockable in C_2 .*

PROOF. By induction on the length of pt . The base case where $|pt| = 0$ is trivial. If $|pt| > 0$ then let $pt = pt' \cdot a$. The induction hypothesis gives that pt' is unblockable in C_1 and $\{\alpha | P_2 \mid \alpha \in \text{caused}(pt')\}$ are unblockable in $\mathcal{T}_p(C_2)$ iff pt' is unblockable in $C_1 \ll C_2$.

First, we prove the \Leftarrow direction by contradiction. That is pt is blockable in in $C_1 \ll C_2$. Then there exists $(s_0^1, s_0^2) \xrightarrow{tr} (r_1', r_2')$ such that $a \notin \text{enabled}((r_1', r_2'))$. By Lemma 4, this is a contradiction to the statements of the left side.

Second, we prove the \Rightarrow direction also by contradiction. If pt is blockable in C_1 , so there exists $s_0^1 \xrightarrow{tr'} r_1'$ with $\pi_1(tr') = pt'$ such that $a \notin \text{enabled}(r_1')$. By induction hypothesis, we know pt' is unblockable in $C_1 \ll C_2$. By Lemma 4, there exists $(s_0^1, s_0^2) \xrightarrow{tr} (r_1', r_2')$ where $\pi_1(tr) = pt'$, but $a \notin \text{enabled}(r_1', r_2')$ because $a \notin \text{enabled}(r_1')$ and $P_2 \subseteq R_1$. This contradicts the fact that $pt' \cdot a$ is unblockable in $C_1 \ll C_2$. If $\{\alpha | P_2 \mid \alpha \in \text{caused}(pt)\}$ are blockable in C_2 , then there exists $r_2 \in S_2$ such that C_2 cannot accept $\text{caused}(r_1, a)$ starting from state r_2 , because $\{\alpha | P_2 \mid \alpha \in \text{caused}(pt')\}$ are unblockable by induction hypothesis. By the transition rules of the automaton $C_1 \ll C_2$, $a \notin \text{enabled}((r_1, r_2))$. This contradicts that pt is unblockable in $C_1 \ll C_2$, and the proof of this direction is also established. \square

3.2 Composition of Interface Models

Every component interface automaton is also a component automaton, so the definitions of product and plugging can also be used to define the composition of interface models. But as shown in Example 3, the product construction on interface models may drop some unblockable traces that could be present in the product of their execution models. So for the interface models, we can compose them only when the plugging condition is satisfied. In that case, the composition of interface models is significant in the sense that unblockable traces are preserved.

However, as demonstrated in Fig. 4, the result is not necessarily a component interface automaton. The figure shows that plugging of two interface automata (the left one is plugged by the right one in the top of Fig. 4) may contain blockable executions, as ac is blockable from the state 1A even though the two interface automata above do not contain any blockable traces.

Hence we define the composition of two pluggable component interface automata as the product construction introduced in the previous section followed by the call of the algorithm \mathcal{I} for removing the blockable executions.

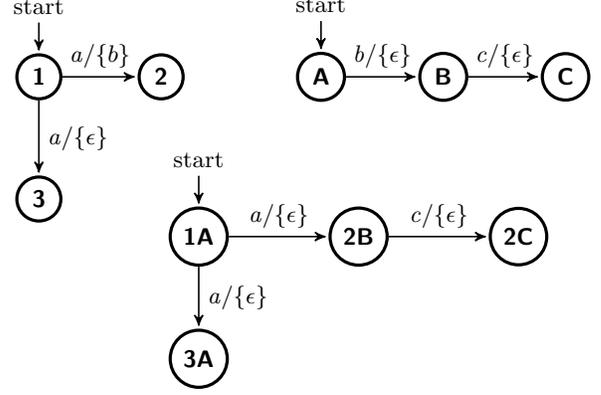


Figure 4: Plugging of interface automata

DEFINITION 8. *Given two component interface automata I_1 and I_2 , if I_1 can be plugged by I_2 , then their composition (plugging), denoted by $I_1 \ll_{\mathcal{I}} I_2$, is defined as $\mathcal{I}(I_1 \ll I_2)$.*

We can now show that plugging of two components is equivalent (w.r.t. to unblockable equivalence) with the composition of the corresponding interface models. This property shows that unblockable provided traces allowed in the composition of interface models are also available in the plugging of their execution models.

THEOREM 4. *Let C_1 be pluggable by C_2 such that $P_2 \subseteq R_1$. Then $\mathcal{I}(C_1) \ll_{\mathcal{I}} \mathcal{I}(C_2) \equiv_{pt} C_1 \ll C_2$.*

PROOF. Theorem 1 shows that $\mathcal{I}(C_1) \equiv_{pt} C_1$, $\mathcal{I}(C_2) \equiv_{pt} C_2$, and $\mathcal{I}(C_1) \ll_{\mathcal{I}} \mathcal{I}(C_2) \equiv_{pt} \mathcal{I}(C_1 \ll C_2)$. From Theorem 3 we get $C_1 \equiv_{pt} C_1 \ll C_2$ and $\mathcal{I}(C_1) \equiv_{pt} \mathcal{I}(C_1) \ll \mathcal{I}(C_2)$. All together we get $\mathcal{I}(C_1) \ll_{\mathcal{I}} \mathcal{I}(C_2) \equiv_{pt} C_1 \ll C_2$. \square

4. COORDINATION OF COMPONENTS

The components discussed in the previous section provide a set of services while causing sequences of required services. A component is like a container of services. For instance, the two components C_{ic} and C_{ip} discussed above interact with other components by providing and invoking services. However, if we want to trigger some services internally, the components will not be enough. For example, after *login* is invoked, we want the status information of the internet component to be shown regularly and automatically. In component-based design, we also need particular processes to interconnect components and such processes coordinate ordinary components by triggering invocations to services provided by these components. Thus, processes can be taken as components that provide nothing but require services according to their control flow.

Invocations to services provided by other components are modeled as *active events*, which means that these events are started by the processes actively. The behavior of a process can be modeled by a traditional finite automata [HMU79], where the alphabet is the set of active events in the process.

DEFINITION 9. *A process automaton is defined as a tuple $A = (S, s_0, E, \delta)$ where*

- $\diamond S$ is a finite set of states,

- ◇ $s_0 \in S$ is the initial state,
- ◇ E is a finite set of active events, and
- ◇ $\delta \subseteq S \times E \times S$ is a transition relation.

As in the component automaton, we write $s \xrightarrow{a} s'$ instead of $(s, a, s') \in \delta$. An alternating sequence of states and active events $exec_p = \langle s_0, a_0, \dots, a_k, s_{k+1} \rangle$ is called an *execution* of the process A if $s_i \xrightarrow{a_i} s_{i+1}$ for all $0 \leq i \leq k$ where s_0 is the initial state. A sequence of active events $\beta = \langle a_0, a_1, \dots, a_k \rangle$ is called a *trace* if there is an execution $exec_p$ such that $exec_p|E = \beta$. The set of traces in A is denoted as $\mathcal{T}_a(A)$.

Now, we will introduce coordination of components by processes. The intuitive idea is to synchronize provided services of the component and the invocation to these services by the process, and the corresponding causality traces will be caused by previous provided service. It is a little different from the product of components. For example, assume for a component automaton C and a process A , there are transitions

$$r \xrightarrow{a/T_1} r' \xrightarrow{b/T_2} r'' \text{ and } s \xrightarrow{b} s',$$

in C and A respectively, then the synchronization will be

$$(r, s) \xrightarrow{a/T_1 \circ T_2} (r'', s'),$$

provided there is no more synchronization transition subsequent to r'' and s' . And if there is one more transition from r' , e.g., $r' \xrightarrow{c/T_3} r_1$ such that c is not required by A , the provided service c should also be available to the environment, which will form additional transitions

$$(r, s) \xrightarrow{a/T_1} (r', s) \text{ and } (r', s) \xrightarrow{c/T_3} (r_1, s).$$

For convenience, we use *conc* for the concatenation of sets of sequences, that is, $conc(\langle T_0, \dots, T_k \rangle) = T_0 \circ \dots \circ T_k$.

DEFINITION 10. A coordination of a component automaton $C = (S_1, s_0^1, P, R, \delta_1)$ by a given process automaton $A = (S_2, s_0^2, E, \delta_2)$ is a component $C \times A = (S, s_0, P', R, \delta)$ where

- ◇ $S = S_1 \times S_2$,
- ◇ $s_0 = (s_0^1, s_0^2)$,
- ◇ $P' = P \setminus E$,
- ◇ δ is the smallest relation constructed as follows.

Let $r_1 \xrightarrow{a/T} r_1' \in \delta_1$ with $a \in P'$ and $r_2 \in S_2$. If

$$r_1' \xrightarrow{tr} r_1'' \text{ and } r_2 \xrightarrow{\pi_1(tr)} r_2'$$

with $enabled(r_1'') \cap enabled(r_2') = \emptyset$ or $enabled(r_1'') \cap P' \neq \emptyset$, then we add

$$(r_1, r_2) \xrightarrow{a/T'} (r_1'', r_2')$$

where

$$T' = T \circ conc(\pi_2(tr)).$$

Given a component automaton $C = (S_1, s_0^1, P, R, \delta_1)$, a process automaton $A = (S_2, s_0^2, E, \delta_2)$, and the coordination $C \times A = (S, s_0, P', R, \delta)$, the next theorem shows the relation between transition relations between the coordination and the component.

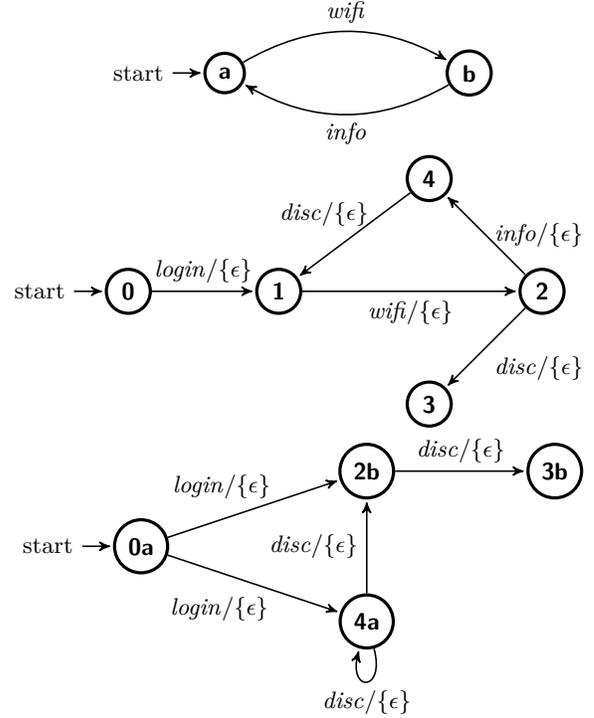


Figure 5: Process wifi-connection A top, component $C_{ic} \otimes C_{ip}$ middle, and Coordination $C_{ic} \times A$ bottom

THEOREM 5. For any execution $(s_0^1, s_0^2) \xrightarrow{tr} (r_1, r_2)$ in δ , we have

$$s_0^1 \xrightarrow{tr_1} r_1 \text{ and } s_0^2 \xrightarrow{tr_2} r_2$$

in C and A respectively, such that

1. $\pi_1(tr) = \pi_1(tr_1)|P'$, that is, there exists an equal provided trace in C when projecting to P' ;
2. $conc(\pi_2(tr)) = conc(\pi_2(tr_1))$, that is, the caused traces in C and $C \times A$ are equal;
3. $tr_2 = \pi_1(tr_1)|E$, showing that the internally triggered sequence of provided events is invoked by the process.

PROOF. For one transition step the proposition can be directly derived from the transition rules of the coordination. The general case can be proven by induction over the length of tr . \square

EXAMPLE 4. A wifi-connection process A is depicted in Fig. 5. The process coordinates internet-connection component by repeatedly requiring wifi. The coordination of this process and internet-connection component $C_{ic} \times A$ shown at the bottom of the figure. For simplicity, we do not draw the event read in the figure. When login is invoked by the environment, wifi is triggered internally and the component causes *unu1* or *unu2* automatically.

5. CONCLUSION AND FUTURE WORK

We have presented a formal model for component interfaces that guarantees unblockable composition of software components. To this end, we have provided an algorithm

to compute the unblockable interface from a component's execution model. This interface serves as a behavioral contract for the component: as long as the environment calls only the sequences of provided services that are specified by the interface and all required services are provided, then non-blocking behavior is guaranteed. Moreover, the complementing operation of coordination of components has been defined so that processes can be used to glue components together in order to support component-based design by building-up new components from the existing ones.

During the construction of a software system from components, one is at the end interested in closed components where all needed required services are resolved, i.e., there are no dangling references to required services. Therefore, we introduced the operators of composition and plugging. We showed that plugging preserves unblockable behavior such that the unblockable sequences of a plugging can be computed from the interfaces of the plugged components. On the other hand, such a result does not hold for the more general composition. The reason is a cyclic dependence between components that may in principle resolve a blockable behavior such that the composition has a larger interface than the one computed from the interfaces of the components.

Due to this restriction, the presented interface model is not suited to reflect low-level interaction protocols with cyclic behaviour, but instead describes, on a higher level, the interfaces of components in service-oriented computing where the non-cyclic service call does usually not impose any practical restriction as mentioned in e.g. [BDF05, CLL⁺07, BBNL08].

There are several open problems left for the future work. We have shown that the plugging operator preserves the unblocking behavior of its constituents, but it is so far not clear for which other composition operators this property holds too. Similarly, it has to be investigated whether the composition operators preserve a refinement relation on components. Intuitively, a refined component provides more services while relying on fewer invocations to required services. It seems that the notion of modal refinement [LT88, LNW07b] can be applicable in this situation. Another research direction are interface models with timing characteristics to support timing analysis of components and scheduling analysis of application processes as well as unblockable behavior in the presence of timed synchronization.

6. REFERENCES

- [AG97] R. Allen and D. Garlan. A formal basis for architectural connection. *ACM Transactions on Software Engineering and Methodology*, 6(3):213–249, 1997.
- [Arb04] F. Arbab. Reo: a channel-based coordination model for component composition. *Mathematical Structures in Computer Science*, 14:329–366, 2004.
- [BBNL08] M. Boreale, R. Bruni, R. D. Nicola, and M. Loreti. Sessions and pipelines for structured service programming. In G. Barthe and F. S. de Boer, editors, *FMOODS*, volume 5051 of *LNCS*, pages 19–38. Springer, 2008.
- [BDF05] M. Bartoletti, P. Degano, and G. L. Ferrari. Enforcing secure service composition. In *CSFW*, pages 211–223. IEEE Computer Society, 2005.
- [CLL⁺07] Z. Chen, X. Li, Z. Liu, V. Stolz, and L. Yang. Harnessing rCOS for tool support: the CoCoME experience. In *Formal methods and hybrid real-time systems*, pages 83–114. Springer, 2007.
- [DAH01a] L. De Alfaro and T. Henzinger. Interface automata. *ACM SIGSOFT Software Engineering Notes*, 26(5):109–120, 2001.
- [dAH01b] L. de Alfaro and T. Henzinger. Interface theories for component-based design. In T. Henzinger and C. Kirsch, editors, *Embedded Software*, volume 2211 of *LNCS*, pages 148–165. Springer, 2001.
- [DAH05] L. De Alfaro and T. Henzinger. Interface-based design. *Engineering Theories of Software-intensive Systems*, 195:83–104, 2005.
- [DLL⁺10] A. David, K. Larsen, A. Legay, U. Nyman, and A. Wasowski. Timed I/O automata: a complete specification theory for real-time systems. In *HSCC*, pages 91–100. ACM, 2010.
- [HMU79] J. Hopcroft, R. Motwani, and J. Ullman. *Introduction to Automata Theory, Languages, and Computation*, volume 2. Addison-Wesley, 1979.
- [LMS09] Z. Liu, C. Morisset, and V. Stolz. rCOS: Theory and tool for component-based model driven development. In F. Arbab and M. Sirjani, editors, *FSEN*, volume 5961 of *LNCS*, pages 62–80. Springer, 2009.
- [LNW07a] K. G. Larsen, U. Nyman, and A. Wasowski. Modal I/O automata for interface and product line theories. In R. D. Nicola, editor, *ESOP*, volume 4421 of *LNCS*, pages 64–79. Springer, 2007.
- [LNW07b] K. G. Larsen, U. Nyman, and A. Wasowski. On modal refinement and consistency. In L. Caires and V. T. Vasconcelos, editors, *CONCUR*, volume 4703 of *LNCS*, pages 105–119. Springer, 2007.
- [LT87] N. A. Lynch and M. R. Tuttle. Hierarchical correctness proofs for distributed algorithms. In *PODC*, pages 137–151, 1987.
- [LT88] K. G. Larsen and B. Thomsen. A modal process logic. In *LICS*, pages 203–210. IEEE Computer Society, 1988.
- [LT89] N. A. Lynch and M. R. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2(3):219–246, 1989.
- [Szy97] C. Szyperski. *Component Software: Beyond Object-Oriented Programming*. Addison-Wesley, 1997.