

Relating two automata-based models of orchestration and choreography [☆]



D. Basile ^{a,*}, P. Degano ^{a,*}, G.L. Ferrari ^{a,*}, E. Tuosto ^{b,*}

^a Dipartimento di Informatica, Università di Pisa, Italy

^b Computer Science Department, University of Leicester, United Kingdom

ARTICLE INFO

Article history:

Received 15 October 2014

Received in revised form 16 September 2015

Accepted 24 September 2015

Available online 12 October 2015

ABSTRACT

We investigate the relations between two automata-based models for describing and studying distributed services, called contract automata and communicating machines. In the first model, distributed services are abstracted away as automata – oblivious of their partners – that coordinate with each other through an orchestrator. The second one is concerned with the interactions occurring between distributed services, that are represented by channel-based asynchronous communications; then services are coordinated through choreography.

We define a notion of strong agreement on contract automata; exhibit a natural mapping from this model to communicating machines with a synchronous semantics; and give conditions to ensure that strong agreement corresponds to well-formed choreography. Then these results are extended to a more liberal notion of agreement and to a fully asynchronous semantics of communicating machines.

© 2015 Elsevier Inc. All rights reserved.

1. Introduction

Among the most popular approaches to the design of distributed coordination, *orchestration* and *choreography* aim to describe the distributed workflow (in other words, how control and data exchanges are coordinated in distributed applications or systems). Intuitively, orchestration yields the description of a distributed workflow from “one party’s perspective” [1], whereas choreography describes the behaviour of involved parties from a “global viewpoint” [2]. In an orchestrated model, the distributed computational components coordinate with each other by interacting with a special component, *the orchestrator*, which at run time dictates how the computation evolves. In a choreographed model, the distributed components autonomously execute and interact with each other on the basis of a local control flow expected to comply with their role as specified in the “global viewpoint”. Here we investigate the relations between two models of distributed coordination: *contract automata* [3] and *communicating machines* [4].

The first model has been recently introduced as a *contract-based* coordination framework where contracts specify the expected behaviour of distributed components oblivious of their communicating partners. The underlying coordination mechanism of contract automata is *orchestration*. In this model a component, or *principal*, is assumed to communicate

[☆] This work has been partially supported by the MIUR project *Security Horizons* and IST-FP7-FET open-IP project *ASCENS*, by the European Union Seventh Framework Programme under grant agreement No. 295261 *Mobility between Europe and Argentina applying Logics to Systems* (MEALS), and by the COST Action IC1201: *Behavioural Types for Reliable Large-Scale Software Systems* (BETTY).

* Corresponding authors.

E-mail addresses: basile@di.unipi.it (D. Basile), degano@di.unipi.it (P. Degano), giangi@di.unipi.it (G.L. Ferrari), emilio@le.ac.uk (E. Tuosto).

messages on some ports according to an automaton specifying its behavioural contract. These messages have to be thought of as directed to an orchestrator synthesised out of the components; the orchestrator directs the interactions in such a way that only executions that “are in agreement” actually happen. In this way, it is possible to transfer the approach of [5,6] to contract automata so to identify misbehaviour of components that do not realise their contract.

We illustrate this with the following simple example. Alice is willing to lend her aeroplane toy, Bob offers a bike toy in order to play with an aeroplane toy, while Carol wants to play with an aeroplane or a bike toy. Let \bar{a} and \bar{b} denote respectively the actions of offering an aeroplane or a bike toy and, dually, a and b denote the corresponding request actions. The contract automata for the principals Alice, Bob, and Carol correspond to the following regular expressions, used for conciseness from here onwards:

$$\text{Alice} = \bar{a} \quad \text{Bob} = \bar{b}.a + a.\bar{b} \quad \text{Carol} = a + b$$

If Alice exchanges her toy with Bob, then all contracts are fulfilled. Instead, if not coordinated, Alice, Bob, and Carol may share their toys in a way that does not fulfil their contracts. Indeed, Alice can give her aeroplane to Bob or Carol, while Carol can receive the bike from Alice or Bob. Now, if Alice gives her aeroplane to Carol the contracts of Alice and Carol are fulfilled, while Bob's is not. In the model of contract automata the coordinator acts as the mother of the three kids who takes their desires and suggests how to satisfy them (and reproaches those who do not act according to their declared contract).

The other model we consider here are *communicating machines*, a sort of automata that can exchange messages through asynchronous (FIFO) buffers. A set of machines constitutes a communicating system. This model has been introduced to study distributed communication protocols and to ensure the correctness of distributed components. Unlike contract automata, communicating machines require no orchestrator since they directly interact with each other through the buffers. Recently, a relation between communicating machines and distributed choreographies has been proved in [7].

We show that the above two models are related, in spite of the different problems they address and the different mechanisms they use for coordination. For this purpose, we first introduce the notion of *strong agreement* on contract automata, that requires the fulfilment of *all* offers and requests, differently than previously introduced notions [3], that we will deal with later on. We then introduce *strongly safe* contract automata, that only accept computations that are in strong agreement.

We consider *convergent* communicating systems, that is those exhibiting successfully terminating computations only. The machines of convergent systems respect their contracts, namely they accomplish their tasks and receive what they look for. Safety and convergence are key notions for semantically linking contract automata and communicating machines. We proceed as follows. We first define a mapping for translating (each principal of) a contract automaton in the corresponding (machine composing a) communication system, and we prove that the computations of the two correspond in a precise way. In the beginning we endow communicating systems with a synchronous semantics, while the general case is considered in Section 6.2.

Then, if a contract automaton is strongly safe the corresponding communicating system is convergent, and vice versa. It should be noted that a given contract automaton can always be transformed into one that is strongly safe, called its (most permissive) *controller*, by adapting standard constructions from the supervisory control theory for discrete event systems [8], as done in [3].

To be more precise, we first establish the above semantic connection by requiring the buffers of communicating machines to contain at most one message and contract automata to well-behave on branching constructs (a notion made precise below).

Later on, we recall the notion of *agreement* on contract automata [3], that only requires the fulfilment of *all* requests. In a sense, this notion allows a contract automaton to be compositionally placed in an unknown environment that may accept the unmatched offers of the automaton. Again, contract automata in agreement and well-behaving on branches are in bijection with the corresponding convergent communicating systems.

Finally, we consider the fully asynchronous semantics of communicating machines, and we prove a weaker result: a strongly safe contract automaton satisfying tighter constraints on branches represents a convergent communicating system.

This paper extends our contribution in [9]. Besides Section 2 on related work, the main new parts concern the notion of agreement just mentioned above, and the new correspondence results for communicating machines with the general semantics, both included in Section 6. Although most of the technicalities on contract automata are inherited from [3], the notion of strong agreement has been introduced here. Finally, we have also revised and simplified some proofs.

Structure of the paper. Section 2 briefly surveys the related work closer to ours. We recall contract automata and communicating finite-state machines in Section 3. Section 4 introduces our new notion of strong agreement on contract automata and the way their controllers are derived. The translation of contract automata into communicating machines is given in Section 5, where we also prove our main theorem of correspondence. In Section 6 we extend our results by relaxing the constraints put on both kinds of automata, i.e. we consider agreement on contract automata and the fully asynchronous semantics for communicating machines. Section 7 works out in full detail an example. Finally, Section 8 concludes, and discusses possible extensions of our results.

2. On choreography and orchestration: related work

To better cast our results within the literature, we first clarify what we mean by choreography and orchestration since, to the best of our knowledge, precise definitions of those concepts are still missing and only intuitive descriptions have been given so far.

There is common consensus that the distinguishing element of a choreographic model is the specification of a so-called *global viewpoint* detailing the interactions among distributed participants and offering a “contract” about their expected communication behaviour in terms of message exchanges. This intuition is best described in W3C words [2]:

Using the Web Services Choreography specification, a contract containing a global definition of the common ordering conditions and constraints under which messages are exchanged, is produced that describes, from a global viewpoint [...] observable behaviour [...]. Each party can then use the global definition to build and test solutions that conform to it. The global specification is in turn realised by combination of the resulting local systems [...]

Noteworthy, the excerpt above points out that local behaviour should then be realised by conforming to the global viewpoint in a “top-down” fashion. Hence, the relations among the global and local specifications are paramount. These aspects are addressed in [10] through an analysis of the relations between the *interaction-oriented* choreographies (i.e., global specifications expressed as interactions) and the *process-oriented* ones (i.e., the local specifications expressed as process algebra terms). A different “bottom-up” approach has been recently introduced in [11] that synthesise choreographies from local specifications. This makes choreography models more flexible (for instance, choreographies have been exploited in [12] as a contract model for service composition). Note that – adapting the terminology of [10] – we use communicating machines as *automata-oriented* choreography, as in [11].

The concept of orchestration is more controversial. In this paper we adopt a widely accepted notion of orchestration [13–15] envisaging the distributed coordination of services as mediated by a distinguished participant that – besides acting as provider of some functionalities – regulates the control flow by exchanging messages with partner services according to their exposed communication interface. In Peltz’s words [1]:

Orchestration refers to an executable business process that can interact with both internal and external Web services. The interactions occur at the message level. They include business logic and task execution order, and they can span applications and organisations to define a long-lived, transactional, multi-step process model. [...] Orchestration always represents control from one party’s perspective.

The “executable process” mentioned by Peltz is called *orchestrator* and specifies the workflow from the “one party’s perspective” describing the interactions with other available services, so to yield a new composed service. This description accounts for a composition model enabling developers to combine existing and independently developed services. The orchestrator then “glues” them together in order to realise a new service, as done for instance in Orc [16]. This is a remarkable aspect since the services combined by an orchestrator are not supposed to have been specifically designed for the service provided by the orchestrator and can in fact be (re)used by other orchestrators for realising different services. Notice that this approach differs from the “bottom-up” one of [11], because synthesised choreographies do not correspond to executable orchestrators.

Other authors consider orchestration as the description of message exchanges among participants from the single participants’ viewpoint *without assuming the presence of an orchestrator*. For instance, in [17,18] the local specifications of a choreography are considered the orchestration model of the choreography itself. We believe that this acceptance is too lax because any distributed application consists of parties that exchange information at will (no matter if realised with channel communication, remote method invocation, etc.). Considering each local specification of a choreography as an orchestration may obscure the matter; rather local specifications are tailored to (and dependent of) the corresponding party of the choreography instead of being independently designed. In other words, for us such local specifications correspond to automata-oriented choreography adopted here and in [11], as well as to the process-oriented choreography of [10]. Instead, in an orchestration model, each participant defines and exposes its own communication pattern which is then (somehow) assembled in an orchestration. We model orchestrators as contract automata and, abstracting from technological aspects, we can describe our approach using Ross-Talbot’s words [19]:

In the case of orchestration we use a service to broker the interactions between the services including the human agents and in the case of choreography we define the expected observable interactions between the services as peers as opposed to mandating any form of brokering.

The dichotomy orchestration–choreography has been discussed in several papers (see e.g., [1]). The only formal results we are aware of that tightly link a choreography to an orchestration framework are in [20]; this paper uses a bisimulation to establish a conformance relation (or its absence) between choreographed and orchestrated computations. We note that the orchestration model of [20] – unlike ours – envisages systems as the parallel composition of orchestrators.

We instead devise conditions to correlate orchestrated computations with local specifications of choreographies so to ensure that the former well-behave and correspond to communication-safe choreographies. (We remark that the “liberal” execution of orchestrated computations not enjoying our conditions would in general be non-compliant with the corresponding choreographic model.) A practical outcome of our results is that contract automata enjoying (strong) agreement can execute without controller, if they are trusted. This yields the further advantage that contract automata are translated into communicating machines that run without any central control: disposing the orchestrator reduces the communication overhead.

The idea behind our notions of agreement is to require, in a composition of participants, the existence of possible “good” executions together with “bad” computations, that the orchestrator eventually cuts off, while the literature has also notions that require composition of participants to have “good” computations, only. Further research problems have been identified and investigated for distributed choreographies, for instance realizability [21,22,17,23], conformance [24–27], or enforcement [28]. Our work departs from the existing literature in that we does not require to start from a global description of the interactions. In fact, we simply assume that each service specifies the interactions it is involved in (oblivious of other partners). Then, we identify (i) the conditions to allow the coordination of a set of services to satisfy each local requirement, and (ii) the conditions for removing the central orchestrator.

3. Background

This section summarises the automata models we use in the paper. Both models envisage distributed computations as enacted by components that interact by exchanging messages. As we will see, in both cases components, abstracted away as automata, yield systems also formalised as automata.

3.1. Contract automata

Before recalling contract automata [3], we fix our notations and preliminary definitions.

Given a set X , as usual, $X^* \stackrel{\text{def}}{=} \bigcup_{n \geq 0} X^n$ is the set of finite words on X (ε is the empty word, ww' is the concatenation of words $w, w' \in X^*$, $w_{(i)}$ denotes the i -th symbol of w , and $|w|$ is the length of w); write x^n for the word obtained by n concatenations of $x \in X$ and x^* for a finite and arbitrarily long repetition of $x \in X$. It will also be useful to consider X^n as a set of tuples and let \vec{x} to range over it. Sometimes, overloading notation (and terminology), we confound tuples on X with words on X (e.g., if $\vec{w} \in X^n$, then $|\vec{w}| = n$ is the length of w and $\vec{w}_{(i)}$ denotes the i -th element of w).

A contract automaton (cf. Definition 2 below) represents the behaviour of a set of principals (possibly a singleton) capable of performing some *actions*; more precisely, as formalised in Definition 1, the actions of contract automata allow them to “advertise” offers, “make” requests, or “handshake” on simultaneous offer/request actions. Consequently, transitions of contract automata will be labelled with tuples of elements in the set $\mathbb{L} \stackrel{\text{def}}{=} \mathbb{R} \cup \mathbb{O} \cup \{\square\}$ where

- requests of principals will be built out of \mathbb{R} while their *offers* will be built out of \mathbb{O} ,
- $\mathbb{R} \cap \mathbb{O} = \emptyset$, and
- $\square \notin \mathbb{R} \cup \mathbb{O}$ is a distinguished label to represent components that stay idle.

We let a, b, c, \dots range over \mathbb{L} and fix an involution $\bar{\cdot} : \mathbb{L} \rightarrow \mathbb{L}$ such that

$$\bar{\mathbb{R}} \subseteq \mathbb{O}, \quad \bar{\mathbb{O}} \subseteq \mathbb{R}, \quad \forall a \in \mathbb{R} \cup \mathbb{O} : \bar{\bar{a}} = a, \quad \text{and} \quad \bar{\square} = \square$$

Definition 1 (Actions). A tuple \vec{a} on \mathbb{L} is

- a *request (action)* on b if and only if \vec{a} is of the form $\square^* b \square^*$ with $b \in \mathbb{R}$
- an *offer (action)* on b if and only if \vec{a} is of the form $\square^* b \square^*$ with $b \in \mathbb{O}$
- a *match (action)* on b if and only if \vec{a} is of the form $\square^* b \square^* \bar{b} \square^*$ with $b \in \mathbb{R} \cup \mathbb{O}$.

We define the relation $\bowtie \subseteq \mathbb{L}^* \times \mathbb{L}^*$ as the symmetric closure of $\dot{\bowtie} \subseteq \mathbb{L}^* \times \mathbb{L}^*$ where $\vec{a}_1 \dot{\bowtie} \vec{a}_2$ if and only if

- \vec{a}_1 and \vec{a}_2 are actions of the same length and
- $\exists b \in \mathbb{R} \cup \mathbb{O} : \vec{a}_1$ is an offer on $b \implies \vec{a}_2$ is a request on b and
- $\exists b \in \mathbb{R} \cup \mathbb{O} : \vec{a}_1$ is a request on $b \implies \vec{a}_2$ is an offer on b .

We write $\vec{a}_1 \bowtie_b \vec{a}_2$ when there is $b \in \mathbb{R} \cup \mathbb{O}$ such that \vec{a}_1 and \vec{a}_2 are actions on b and $\vec{a}_1 \dot{\bowtie} \vec{a}_2$.

We are ready to introduce the notion of a contract automaton made of n principal, each of which represents an entity involved in the service.

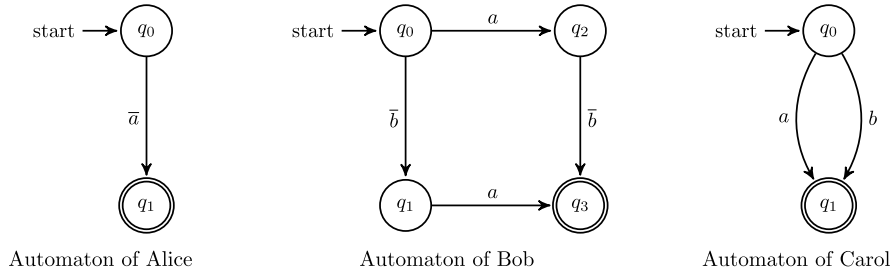


Fig. 1. The contract automata of Alice, Bob and Carol.

Definition 2 (Contract Automata). Let \mathcal{Q} (ranged over by q_1, q_2, \dots) be a finite set of states. A *contract automaton of rank n* is a (finite-state) automaton $\mathcal{A} = \langle \mathcal{Q}^n, \vec{q}_0, \mathbb{L}^n, T, F \rangle$, where

- $\vec{q}_0 \in \mathcal{Q}^n$ is the *initial state*
- $F \subseteq \mathcal{Q}^n$ is the set of *accepting states*
- $T \subseteq \mathcal{Q}^n \times \mathbb{L}^n \times \mathcal{Q}^n$ is the set of *transitions* such that $(\vec{q}, \vec{a}, \vec{q}') \in T$ if and only if
 - if $\vec{a}_{(i)} = \square$ then $\vec{q}_{(i)} = \vec{q}'_{(i)}$ (i.e., the i -th principal stays idle) and
 - \vec{a} is either a request, or an offer, or else a match action

A *principal* is a contract automaton \mathcal{A} of rank 1 such that, for any two transitions (q_1, a_1, q'_1) and (q_2, a_2, q'_2) in \mathcal{A} , it is not the case that $a_1 \bowtie a_2$.

Note that a principal is not allowed to require what it offers.

Example 1. Consider again the simple example of the Introduction. The principal automata of Alice, Bob, and Carol are in Fig. 1, with the usual graphical conventions.

Given a contract automaton of rank n $\mathcal{A} = \langle \mathcal{Q}^n, \vec{q}_0, \mathbb{L}^n, T, F \rangle$, the standard definitions and constructions of finite-state automata apply. In particular,

- the configurations of \mathcal{A} are pairs in $\mathcal{Q}^n \times (\mathbb{L}^n)^*$ of strings of n -tuples of labels and states of \mathcal{A} ;
- \mathcal{A} moves from (\vec{q}, w) to (\vec{q}', w') , written $(\vec{q}, w) \xrightarrow{\vec{a}} (\vec{q}', w')$, if and only if $w = \vec{a}w'$ and $(\vec{q}, \vec{a}, \vec{q}') \in T$; we write $(\vec{q}, w) \rightarrow (\vec{q}', w')$ when \vec{a} is immaterial and $\vec{q} \xrightarrow{\vec{a}} \vec{q}'$ when w is immaterial (note that transitions are triples, instead);
- the *language* of \mathcal{A} is $\mathcal{L}(\mathcal{A}) = \{w \mid (\vec{q}_0, w) \rightarrow^* (\vec{q}, \varepsilon), \vec{q} \in F\}$ where \rightarrow^* is the reflexive and transitive closure of \rightarrow . As usual, $\vec{q}_1 \xrightarrow{\vec{a}_1 \dots \vec{a}_m} \vec{q}_{m+1}$ shortens $\vec{q}_1 \xrightarrow{\vec{a}_1} \vec{q}_2 \dots \vec{q}_m \xrightarrow{\vec{a}_m} \vec{q}_{m+1}$ (for some $\vec{q}_2, \dots, \vec{q}_m$) and we say that \vec{q}_1 is *reachable* in \mathcal{A} if $\vec{q}_0 \xrightarrow{w} \vec{q}_1$; finally $\vec{q} \not\rightarrow$ if and only if for no \vec{q}' it is the case that $\vec{q} \rightarrow \vec{q}'$.

We now borrow from [3] the product operation of contract automata, that is similar to the one introduced in [29]. Given a finite set of contract automata, this operation basically yields the contract automaton that interleaves all their transitions while forcing synchronisations when two contract automata are in states ready to “handshake” (i.e., they can fire complementary request/offer actions).

Definition 3 (Product). For $i \in \{1, \dots, h\}$, let $\mathcal{A}_i = \langle \mathcal{Q}^{n_i}, \vec{q}_{0_i}, \mathbb{L}^{n_i}, T_i, F_i \rangle$ be a contract automaton of rank n_i . The *product* of $\mathcal{A}_1, \dots, \mathcal{A}_h$, denoted as $\bigotimes_{i \in \{1, \dots, h\}} \mathcal{A}_i$, is the contract automaton $\langle \mathcal{Q}^n, \vec{q}_0, \mathbb{L}^n, T, F \rangle$ of rank $n = n_1 + \dots + n_h$ where:

- $\vec{q}_0 = \vec{q}_{0_1} \dots \vec{q}_{0_h}$
- $F = F_1 \times \dots \times F_h$
- T is the smallest subset of $\mathcal{Q}^n \times \mathbb{L}^n \times \mathcal{Q}^n$ such that $(\vec{q}, \vec{c}, \vec{q}') \in T$ if and only if, letting $\vec{q} = \vec{q}_1 \dots \vec{q}_h \in \mathcal{Q}^n$, $\vec{c} = \vec{c}_1, \dots, \vec{c}_h \in \mathbb{L}^n$,


either there are $1 \leq i < j \leq h$ such that $(\vec{q}_i, \vec{a}_i, \vec{q}'_i) \in T_i$, $(\vec{q}_j, \vec{a}_j, \vec{q}'_j) \in T_j$, $\vec{a}_i \bowtie \vec{a}_j$, and

$$\begin{cases} \vec{c}_i = \vec{a}_i, \vec{c}_j = \vec{a}_j, \text{ and } \vec{c}_l = \square^{n_l} \text{ for } l \in \{1, \dots, h\} \setminus \{i, j\} \\ \text{and} \\ \vec{q}' = \vec{q}_1 \dots \vec{q}_{i-1} \vec{q}'_i \vec{q}_{i+1} \dots \vec{q}_{j-1} \vec{q}'_j \vec{q}_{j+1} \dots \vec{q}_h \end{cases}$$

or for each $l \neq i \in \{1, \dots, h\}$, if $(\vec{q}_i, \vec{a}_i, \vec{q}'_i) \in T_i$ then

$$\vec{c}_i = \vec{a}_i, \quad \vec{c}_l = \square^{n_l}, \text{ and } \vec{q}' = \vec{q}_1 \dots \vec{q}_{i-1} \vec{q}'_i \vec{q}_{i+1} \dots \vec{q}_h$$

also, for all $j \neq i$ and $(\vec{q}_j, \vec{a}_j, \vec{q}'_j) \in T_j$ it does not hold that $\vec{a}_i \bowtie \vec{a}_j$.

Fig. 2. Alice \otimes Bob \otimes Carol.

Example 2. The contract automaton in Fig. 2 is the product of the contract automata in Fig. 1. Notice that only match actions leave the states \vec{q}_0 and \vec{q}_3 , indeed principals can handshake and so there are no outgoing offer and request transitions.

Remark 1. The product in Definition 3 is not associative: consider again Example 1. Then the composition the automaton for Carol with the product of those for Alice and Bob is different from the composition of the automaton of Alice with the product of the other two. An alternative definition of associative product is given in [3]. Briefly, before computing the product, all its arguments are projected to their principals. Aiming at a simpler presentation, we opted here for the non-associative product, but we stress that our results hold also with the associative one, because we are composing principals only.

Hereafter, we assume that all contract automata of rank $n > 1$ are the product of n principals. Also, we consider deterministic contract automata only. These assumptions can be relaxed at the cost of adding some technical intricacies.

3.2. Communicating machines

Communicating machines [4] are a simple automata-based model introduced to specify and analyse systems made of agents interacting via asynchronous message passing. We adapt the original definitions and notation from [4] and [30] to our needs; in particular, the only relevant difference with the original model is that we add the set of final states. Let \mathcal{P} be a finite set of *participants* (ranged over by p, q, r, s , etc.) and $C \stackrel{\text{def}}{=} \{pq \mid p, q \in \mathcal{P} \text{ and } p \neq q\}$ be the set of *channels*.

The set of *actions* is $\text{Act} \stackrel{\text{def}}{=} C \times (\mathbb{R} \cup \mathbb{O})$ and it is ranged over by ℓ ; we abbreviate (sr, \bar{a}) with $\bar{a}@_{sr}$ when $\bar{a} \in \mathbb{O}$, representing the action of *sending* a from machine s to r . Similarly, we shorten (sr, a) with $a@_{sr}$ when $a \in \mathbb{R}$, representing the *reception* of a by r .

Definition 4 (CFSM). Given a finite set of states Q , a *communicating finite state machine* is an automaton $M = (Q, q_0, \text{Act}, \delta, F)$ where

- $q_0 \in Q$ is the *initial state*,
- $\delta \subseteq Q \times \text{Act} \times Q$ is the set of *transitions*,
- $F \subseteq Q$ is the set of final, accepting states.

We say that M is *deterministic* if and only if for all states $q \in Q$ and all actions $\ell \in \text{Act}$, if $(q, \ell, q'), (q, \ell, q'') \in \delta$ then $q' = q''$. Finally, we write $\mathcal{L}(M) \subseteq \text{Act}^*$ for the language on Act accepted by the automaton, i.e. the machine M .

The notion of deterministic communicating finite state machines adopted here differs from the standard one (e.g., the one in [30]) which requires that, for any state q , if $(q, \bar{a}@_{sr}, q') \in \delta$ and $(q, \bar{b}@_{sr}, q'') \in \delta$ then $a = b$ and $q' = q''$. We use our variant because it reflects the semantics of contract automata better. Indeed, hereafter, we will only consider deterministic communicating finite state machines.

The communication model of communicating finite state machines (cf. Definitions 5 and 6) is based on (unbounded) FIFO buffers, that actually are the elements in C . They are to be intended as the channels that the participants use to exchange




Fig. 3. Three communicating machines.

messages. To spare another syntactic category and cumbersome definitions, we draw the messages appearing in the buffers of communicating finite state machines from the set of requests \mathbb{R} . Recall that the set of participants \mathcal{P} is finite.

Definition 5 (*Communicating systems*). Given a communicating finite state machine $M_p = (Q_p, q_{0p}, \text{Act}, \delta_p, F_p)$ for each $p \in \mathcal{P}$, the tuple $S = (M_p)_{p \in \mathcal{P}}$ is a *communicating system*, belonging to the set CS.

A *configuration* of S is a pair $s = (\vec{q}; \vec{u})$ where $\vec{q} = (q_p)_{p \in \mathcal{P}}$ with $q_p \in Q_p$ and where $\vec{u} = (u_{pq})_{pq \in \mathcal{C}}$ with $u_{pq} \in \mathbb{R}^*$; the component \vec{q} is the *control state* and $q_p \in Q_p$ is the *local state* of machine M_p , while \vec{u} represents the contents of the (FIFO buffers of the) channels.

The *initial configuration* of S is $s_0 = (\vec{q}_0; \vec{\epsilon})$ with $\vec{q}_0 = (q_{0p})_{p \in \mathcal{P}}$ and $\vec{\epsilon}$ is the vector of empty channels.

Hereafter, we fix a machine $M_p = (Q_p, q_{0p}, \text{Act}, \delta_p, F_p)$ for each participant $p \in \mathcal{P}$ and we let $S = (M_p)_{p \in \mathcal{P}}$ be the corresponding system. The definition below formalises a computation step of a communicating system: if a machine M_s sends a message a to a machine M_r then a is inserted on (the FIFO buffer of) the channel sr connecting the two, rendered by $sr \cdot a$ (condition 1 below). The element is then read from the channel, if this is in the form $a \cdot sr$ being a the top of the channel (condition 2 below). In both cases, no other machine is affected in the step.

Definition 6 (*Reachable state*). A configuration $s' = (\vec{q}'; \vec{u}')$ is *reachable* from another configuration $s = (\vec{q}; \vec{u})$ by firing $\ell \in \text{Act}$, written $s \xrightarrow{\ell} s'$, if there exists $a \in \mathbb{R}$ such that:

1. if $\ell = \bar{a}@sr$ then $(\vec{q}_{(s)}, \ell, \vec{q}'_{(s)}) \in \delta_s$ and for all $p \neq s$, $\vec{q}'_{(p)} = \vec{q}_{(p)} \wedge \vec{u}'_{(sr)} = \vec{u}_{(sr)} \cdot a$;
2. if $\ell = a@sr$ then $(\vec{q}_{(r)}, \ell, \vec{q}'_{(r)}) \in \delta_r$ and for all $p \neq r$, $\vec{q}'_{(p)} = \vec{q}_{(p)} \wedge \vec{u}'_{(sr)} = a \cdot \vec{u}_{(sr)}$;
3. and, in both (1) and (2) above, for all $pq \neq sr$, $\vec{u}'_{(pq)} = \vec{u}_{(pq)}$.

As usual, the computation $s_1 \xrightarrow{\ell_1 \dots \ell_m} s_{m+1}$ shortens $s_1 \xrightarrow{\ell_1} s_2 \dots s_m \xrightarrow{\ell_m} s_{m+1}$ (for some s_2, \dots, s_m).

The set of *reachable configurations* of S is $RS(S) = \{s \mid s_0 \rightarrow^* s\}$.

Example 3. Fig. 3 shows a graphical representation of a communicating system made of three communicating machines, used later on to illustrate the mapping from contract automata to communicating machines.

Similarly to finite state automata, communicating machines have one initial state and a set of final states. A transition has a channel as label, e.g. $\bar{c}@BC$, indicating that the action c is sent by participant B to the buffer of the participant C ; similarly for a receive action. Note that channels of communicating machines specify for each action the sender and the receiver, which is not the case for contract automata.

3.3. Notational synopsis

To avoid cumbersome repetitions, through the paper we assume fixed a contract automaton $\mathcal{A} = \langle \mathcal{Q}^n, \vec{q}_0, \mathbb{L}^n, T, F \rangle$ of rank n . For readability, we summarise the notations introduced so far in Table 1.

Finally, we assume that the states of any automaton/machine are built out of a fixed universe \mathcal{Q} (of states).

4. Enforcing agreement

This section introduces the original notion of *strong agreement* on contract automata, that elaborates the notions of *agreement* and *weak agreement* introduced in [3]. The three notions differ on the conditions for the fulfilment of an interaction between different principals, and share the basic requirement that all the requests of principals are satisfied. Briefly, an *agreement* exists if all the requests, but not necessarily all the offers, are satisfied synchronously. Intuitively, this means that the orchestrator guarantees two principals that their requests are synchronously matched by complementary actions.

Table 1
Notational synopsis.

X^*	set of finite words w on an alphabet X ; ε is the empty word;
$ w , w_{(i)}$	juxtaposition of words is denoted by $_ \cdot _$ the length of w and its i -th symbol, respectively
w^n (w^* , resp.)	w concatenated n -times (arbitrarily many, resp.) with itself
\bar{z} or $(z_i)_{1 \leq i \leq n}$	indexed tuples of states, labels, channel contents, ...
\mathbb{L}	labels
\mathbb{R}	request labels, ranged over by \bar{a}, \bar{b}, \dots
\mathbb{O}	offer labels, ranged over by a, b, \dots
$\square \notin \mathbb{R} \cup \mathbb{O}$	idle label
\mathcal{A}	a contract automaton of rank n the language of which is $\mathcal{L}(\mathcal{A})$
\mathcal{P}	set of participants (ranged over by p, q, i, j, A, B, C , etc.)
C	set of channels, i.e. pairs of participants; (ranged over by pq)
$u \in \mathbb{R}^*$	the FIFO buffer of a channel (top on left, bottom on right)
M_p	communicating machine of participant p
S	a system of communicating machines
$\text{Act} = C \times (\mathbb{R} \cup \mathbb{O})$	actions ℓ , either $\bar{a}@pq$ (sending offer \bar{a}) or $a@pq$ (receiving request a)
$s = (\bar{q}; \bar{u})$	a configuration
$RS(S)$	the set of the reachable configurations of S

Instead, a *weak agreement* exists when request actions can be performed “on credit”. In other words, a computation yields weak agreement when the fulfilment of a request action can happen after the action has been taken. Intuitively, this corresponds to an asynchronous communication admitting inputs taken on credit, provided that obligations will be honoured later on.

Here, we focus on *strongly safe* contract automata, i.e. those enjoying the property of *strong agreement*. This notion is a tighter version of the ones sketched above: it requires *synchronous* fulfilment of all offers and requests. In Section 5 we will show how this condition corresponds to interactions between communicating machines that always succeed.

Definition 7 (Strong Agreement and Strong Safety). A *strong agreement* on \mathbb{L} is a finite sequence of match actions. We let \mathfrak{Z} denote the set of all strong agreements on \mathbb{L} .

A contract automaton \mathcal{A} is *strongly safe* if $\mathcal{L}(\mathcal{A}) \subseteq \mathfrak{Z}$. We say that \mathcal{A} *admits strong agreement* when $\mathcal{L}(\mathcal{A}) \cap \mathfrak{Z} \neq \emptyset$.

We can compose a set of principals so to obtain a strongly safe contract automaton adopting an approach borrowed from the supervisory control theory for discrete event systems [8]. In this theory, discrete event systems are basically automata where accepting states represent successful termination. Instead, *forbidden states* are those that may lead to failures, and thus should not be traversed in “good” computations. The purpose is then to synthesise a controller that enforces a given system to never pass through a forbidden state. The supervisory control theory distinguishes between *controllable* events (those events that the controller can disable) and *uncontrollable* events (those that are always enabled). Moreover, the theory partitions events in *observable* and *unobservable*; the latter being a subset of uncontrollable events. It is known that if all events are observable then a maximally permissive controller exists that never blocks a good computation [8].

We assume that the orchestrator can prescribe any kind of action to be forbidden, so all the events will be controllable. Since the behaviours that we want to enforce in \mathcal{A} are exactly those traces labelled by words in $\mathfrak{Z} \cap \mathcal{L}(\mathcal{A})$, we then specialise the notions of supervisory control theory by defining

- observable events to be all offer, request, and match actions;
- forbidden events to be non-match actions.

Definition 8 (Controller). A (strong) controller of \mathcal{A} is a contract automaton $KS_{\mathcal{A}}$ such that $\mathcal{L}(KS_{\mathcal{A}}) \subseteq \mathfrak{Z} \cap \mathcal{L}(\mathcal{A})$.

The *most permissive (strong) controller* of \mathcal{A} is the controller $KS_{\mathcal{A}}$ such that $\mathcal{L}(KS'_{\mathcal{A}}) \subseteq \mathcal{L}(KS_{\mathcal{A}})$ for all $KS'_{\mathcal{A}}$ controllers of \mathcal{A} .

It is immediate that the most permissive controller is unique up-to language equivalence.

Example 4. The most permissive controller of the contract automaton in Fig. 2 has of the states \bar{q}_0 , \bar{q}_1 , \bar{q}_3 , and \bar{q}_4 with transitions $(\bar{q}_0, (\bar{a}, a, \square), \bar{q}_1)$, $(\bar{q}_3, (\bar{a}, a, \square), \bar{q}_4)$, $(\bar{q}_0, (\square, \bar{b}, b), \bar{q}_3)$, and $(\bar{q}_1, (\square, \bar{b}, b), \bar{q}_4)$. Such controller is easily obtained by applying the construction of Lemma 1 below.




Fig. 4. The controlled system of the contract automaton in Fig. 2.

Proposition 1. If $KS_{\mathcal{A}}$ is the most permissive controller of \mathcal{A} then $\mathcal{L}(KS_{\mathcal{A}}) = \mathfrak{Z} \cap \mathcal{L}(\mathcal{A})$.

Proof. By contradiction, assume $\mathcal{L}(KS_{\mathcal{A}}) \subset \mathfrak{Z} \cap \mathcal{L}(\mathcal{A})$. Since $\mathfrak{Z} \cap \mathcal{L}(\mathcal{A})$ is the intersection of two regular languages and all actions are controllable, there exists a contract automaton $KS'_{\mathcal{A}}$ accepting it (cf. [8]). By definition, $KS'_{\mathcal{A}}$ is a controller of \mathcal{A} strictly containing $\mathcal{L}(KS_{\mathcal{A}})$, contradicting the hypothesis that $\mathcal{L}(KS_{\mathcal{A}})$ is the most permissive controller. \square

It is convenient to introduce the following definition.

Definition 9. A state \vec{q} is *redundant* in the contract automaton \mathcal{A} if, and only if, no accepting state of \mathcal{A} can be reached from \vec{q} .

Lemma 1 (MPC). A contract automaton is the most permissive controller of \mathcal{A} if and only if it is language-equivalent to

$$KS_{\mathcal{A}} \stackrel{\text{def}}{=} \langle \mathcal{Q}^n, \vec{q}_0, \mathbb{L}^n, T_K \setminus \{(\vec{q}, a, \vec{q}') \mid \vec{q} \text{ or } \vec{q}' \text{ is redundant in } K\}, F \rangle$$

where $K = \langle \mathcal{Q}^n, \vec{q}_0, \mathbb{L}^n, T_K = \{t \in T \mid t \text{ is a match transition}\}, F \rangle$ is the sub-automaton of \mathcal{A} the transition relation of which consists of the match transitions of \mathcal{A} only.

Proof. By definition K is a sub-automaton of \mathcal{A} and, by construction, the transitions of $KS_{\mathcal{A}}$ are a subset of the transitions of \mathcal{A} and K contains only match transition, hence $\mathcal{L}(KS_{\mathcal{A}}) \subseteq \mathfrak{Z} \cap \mathcal{L}(\mathcal{A})$. Therefore $KS_{\mathcal{A}}$ is a controller of \mathcal{A} and we have just to prove that $\mathcal{L}(KS_{\mathcal{A}}) = \mathfrak{Z} \cap \mathcal{L}(\mathcal{A})$. We proceed by contradiction.

Let $w \in (\mathfrak{Z} \cap \mathcal{L}(\mathcal{A})) \setminus \mathcal{L}(KS_{\mathcal{A}})$. Since if $\varepsilon \in \mathcal{L}(\mathcal{A})$ then $\varepsilon \in \mathcal{L}(KS_{\mathcal{A}})$, we have $w \neq \varepsilon$ and there must be a transition $t = (\vec{q}, \vec{a}, \vec{q}')$ of \mathcal{A} not in $KS_{\mathcal{A}}$ in the accepting path of w (which is unique since we consider deterministic contract automata only), otherwise $w \in \mathcal{L}(KS_{\mathcal{A}})$. We know that \vec{a} is a match action because $w \in \mathfrak{Z}$, and \vec{q}, \vec{q}' are not redundant states of \mathcal{A} because the transition belongs to an accepting path. Hence there must be $t \in KS_{\mathcal{A}}$, since by construction match transitions between non-redundant states are in $KS_{\mathcal{A}}$. \square

From now onwards, we consider as most permissive controller the one obtained by the construction of Lemma 1. We then define the *controlled system* of a contract automaton \mathcal{A} , using such controller. The purpose of this new automaton is to identify the match transitions of \mathcal{A} together with those that lead “outside” of the controller; in its definition we use the distinguished state $\perp \notin \mathcal{Q}^n$ (for any n).

Definition 10 (Controlled system). Let $KS_{\mathcal{A}} = \langle \mathcal{Q}^n, \vec{q}_0, \mathbb{L}^n, T' \subseteq T, F \rangle$ be the most permissive controller of \mathcal{A} . The *controlled system* of \mathcal{A} under $KS_{\mathcal{A}}$ is the automaton

$$KS_{\mathcal{A}}/\mathcal{A} = \langle \mathcal{Q}^n \cup \{\perp\}, \vec{q}_0, \mathbb{L}^n, T'', F \rangle$$

such that

$$T'' = T' \cup \{(\vec{q}, \vec{a}, \perp) \mid \vec{q} \text{ reachable in } KS_{\mathcal{A}} \text{ and } \exists \vec{q}' \in \mathcal{Q}^n : (\vec{q}, \vec{a}, \vec{q}') \in T \setminus T'\}$$

Example 5. The controlled system of the contract automaton in Fig. 2 is obtained by adding the transitions $(\vec{q}_1, (\square, \square, a), \perp)$, $(\vec{q}_0, (\bar{a}, \square, a), \perp)$ to the most permissive controller of Example 4. It is displayed in Fig. 4.

It is worth remarking that the transitions reaching \perp in the controlled system of \mathcal{A} identify the start of the computations in \mathcal{A} which lead to violations of strong agreement.

In the next definition, we introduce a notion of strong liability, to single out the principals that are potentially responsible of the divergence from the expected behaviour.

Definition 11 (Strong Liability). Given a controlled system $KS_{\mathcal{A}}/\mathcal{A}$ of rank n , the set of *liable* principals on a trace $w \in \mathcal{L}(\mathcal{A})$ is given by:

$$Liable(KS_{\mathcal{A}}/\mathcal{A}, w) = \{i \leq n \mid (\vec{q}_0, w) \rightarrow^* (\vec{q}, \vec{a}w') \rightarrow (\perp, w'') \text{ in } KS_{\mathcal{A}}/\mathcal{A}, \vec{a}_{(i)} \neq \square\}$$

The *potentially liable principals* in $KS_{\mathcal{A}}/\mathcal{A}$ are:

$$Liable(KS_{\mathcal{A}}/\mathcal{A}) \stackrel{\text{def}}{=} \bigcup_{w \in \mathcal{L}(\mathcal{A})} Liable(KS_{\mathcal{A}}/\mathcal{A}, w)$$

Finally, let $TLiable(KS_{\mathcal{A}}/\mathcal{A})$ denote the set of transitions of \mathcal{A} that make principals liable:

$$TLiable(KS_{\mathcal{A}}/\mathcal{A}) \stackrel{\text{def}}{=} \{(\vec{q}, \vec{a}, \perp) \in T_{KS_{\mathcal{A}}/\mathcal{A}} \mid \vec{q} \in \mathcal{Q}^n, \vec{a} \in (\mathbb{L}^*)^n\}$$

Note that the transition labelled by \vec{a} in [Definition 11](#) is the first which diverges from the expected path (since, by [Definition 10](#), state \perp has no outgoing transitions). Indeed a liable principal may fire an action taking the computation away from agreement.

Example 6. The liable indexes of the contract automaton in [Fig. 4](#) are 1 and 3, corresponding to Alice and Carol respectively; the transitions that make them liable are respectively $(\vec{q}_0, (\vec{a}, \square, a), \perp)$ and $(\vec{q}_1, (\square, \square, a), \perp)$. The former liable transition is a match that leads to a non-match transition. So, by inspecting the labels, we easily identify (the indexes of) the liable principals.

5. From contract automata to communicating machines

The translation of a principal into a communicating machine is conceptually straightforward as the two automata are almost isomorphic, apart from their labels. Recall that the principals in a contract automaton can fire transitions not matched by other principals. To account for this kind of “openness,” the [Definition 12](#) below uses the new “ $-$ ” symbol to represent a special, “anonymous” participant, distinguished from those composing the contract automaton in hand, and playing the role of the environment. For this reason, we will assume from now onwards that actions in Act are built on $C \stackrel{\text{def}}{=} \{p\bar{q} \mid p, q \in \mathcal{P} \cup \{-\} \text{ and } p \neq q\}$.

Definition 12 (Translation). For a participant $p \in \mathcal{P}$, let $\llbracket _ \rrbracket_p : \mathbb{L}^n \rightarrow \text{Act}$ be defined as:

$$\llbracket \vec{a} \rrbracket_p = \begin{cases} \vec{a} @ i \bar{j} & \text{if } \vec{a} \text{ is a match action and } i \text{ and } j \text{ are such that} \\ & \vec{a}_{(i)} \in \mathbb{O} \text{ and } \vec{a}_{(j)} \in \mathbb{R} \text{ and } p = i \\ a @ i \bar{j} & \text{if } \vec{a} \text{ is a match action and } i \text{ and } j \text{ are such that} \\ & \vec{a}_{(i)} \in \mathbb{O} \text{ and } \vec{a}_{(j)} \in \mathbb{R} \text{ and } p = j \\ \vec{a} @ i - & \text{if } \vec{a} \text{ is an offer action and } i \text{ is such that } \vec{a}_{(i)} \in \mathbb{O} \text{ and } p = i \\ a @ - \bar{j} & \text{if } \vec{a} \text{ is a request action and } j \text{ is such that } \vec{a}_{(j)} \in \mathbb{R} \text{ and } p = j \\ \varepsilon & \text{otherwise} \end{cases}$$

The translation of \mathcal{A} to a communicating finite state machine is given by the map

$$\llbracket \mathcal{A} \rrbracket_p \stackrel{\text{def}}{=} (\mathcal{Q}, \vec{q}_{0(p)}, \text{Act}, \{(\vec{q}_{(p)}, \llbracket \vec{a} \rrbracket_p, \vec{q}'_{(p)}) \mid (\vec{q}, \vec{a}, \vec{q}') \in T \text{ and } \llbracket \vec{a} \rrbracket_p \neq \varepsilon\}, F)$$

The communicating system *corresponding* to the contract automaton \mathcal{A} is $S(\mathcal{A}) = (\llbracket \mathcal{A} \rrbracket_p)_{p \in \{1, \dots, n\}}$.

The following example illustrates the composition of three contract automata, its most permissive controller, and its translation into a system of three communicating machines.

Example 7. [Fig. 5](#) shows three contract automata A , B , and C (top), their product $A \otimes B \otimes C$ (middle), with initial state $\vec{q}_0 = \langle q_{01}, q_{02}, q_{03} \rangle$, and the corresponding most permissive controller $KS_{A \otimes B \otimes C}$ (bottom). The translations $\llbracket KS_{A \otimes B \otimes C} \rrbracket_A$, $\llbracket KS_{A \otimes B \otimes C} \rrbracket_B$, and $\llbracket KS_{A \otimes B \otimes C} \rrbracket_C$ as per [Definition 12](#) yields the three communicating machines in [Fig. 3](#).

We constrain the behaviour of communicating machines, so to make it easier reflecting in this model the notion of strong agreement defined on contract automata. Intuitively, the new semantics, called *1-buffer semantics*, only allows a machine M_p to send a message \vec{a} to a partner M_q if in the communicating system S all the channels are empty. To specify the *1-buffer semantics*, it suffices to constrain the component $\vec{u}_{(pq)}$ of the reachable configurations $RS(S)$ of S , and only keeping the transitions between the selected configurations. We also define *convergent* communicating systems that always reach a final configuration and so they are *deadlock-free*.

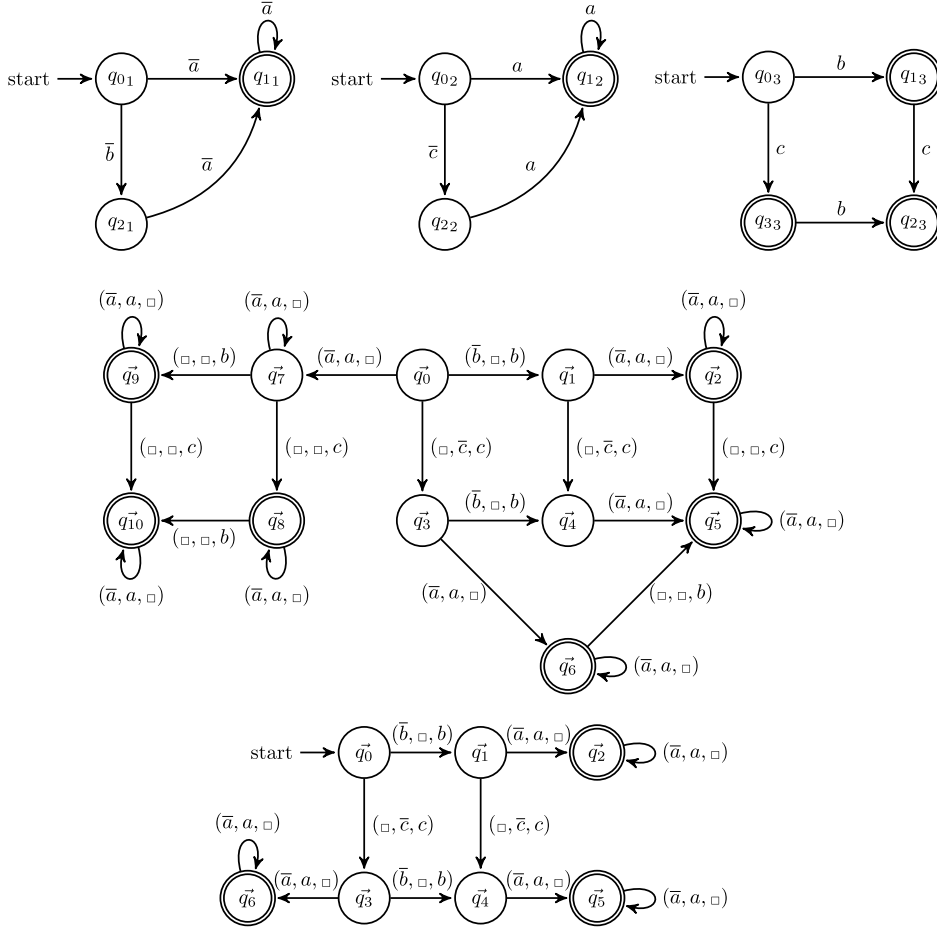


Fig. 5. From left to right, top line: the contract automata A, B, C ; second line: the contract automaton $A \otimes B \otimes C$, the controlled automaton of which $KS_{A \otimes B \otimes C}$ is in the third line; the communicating machines $\llbracket KS_{A \otimes B \otimes C} \rrbracket_A, \llbracket KS_{A \otimes B \otimes C} \rrbracket_B, \llbracket KS_{A \otimes B \otimes C} \rrbracket_C$ corresponding to the contract automata in the first line are shown in Fig. 3.

Definition 13 (1-buffer, convergence, deadlock). Let $S = (M_p)_{p \in \mathcal{P}}$ be a communicating system. A configuration $s = (\vec{q}; \vec{u})$ of S is *stable* if and only if $\vec{u} = \vec{\varepsilon}$; additionally, s is *final* if it is stable and $\vec{q} \in (F_p)_{p \in \mathcal{P}}$.

The transition relation of the 1-buffer semantics of S is the following

$$\rightarrow \stackrel{\text{def}}{=} \rightarrow \cap (RS_{\leq 1}(S) \times \text{Act} \times RS_{\leq 1}(S))$$

where \rightarrow is the relation introduced in Definition 6 and

$$RS_{\leq 1}(S) \stackrel{\text{def}}{=} \{(\vec{q}; \vec{u}) \in RS(S) \mid (\vec{q}; \vec{u}) \text{ is stable or}$$

$$\exists p q \in C : \exists a \in \mathbb{R} : \vec{u}_{(pq)} = a \wedge \forall sr \neq pq. \vec{u}_{(sr)} = \varepsilon\}$$

We say that the system S is *convergent* (with the 1-buffer semantics) if and only if for every reachable configuration $(\vec{q}; \vec{u}) \in RS_{\leq 1}(S)$, there exists a final configuration s such that

$$(\vec{q}; \vec{u}) \rightarrow^* s.$$

Moreover a configuration $(\vec{q}; \vec{u})$ is a *deadlock* if and only if it is not final and $(\vec{q}; \vec{u}) \not\rightarrow$.

In order to relate the computations of contract automata and those of communicating systems, it is convenient to define a translation from the first to the second ones, as follows.

Definition 14. Given a sequence of n -tuples of actions $\varphi \in (\mathbb{L}^n)^*$, we define

$$\llbracket \varphi \rrbracket \stackrel{\text{def}}{=} \begin{cases} \vec{a} @ i j \vec{a} @ i j \llbracket \varphi' \rrbracket & \text{if } \varphi = \vec{a} \varphi' \text{ and } \vec{a} \text{ is a match action on } a \\ & \text{with } \vec{a}_{(i)} \in \mathbb{O} \text{ and } \vec{a}_{(j)} \in \mathbb{R} \\ \vec{a} @ i - \llbracket \varphi' \rrbracket & \text{if } \varphi = \vec{a} \varphi' \text{ and } \vec{a} \text{ is an offer action on } a \\ & \text{with } \vec{a}_{(i)} \in \mathbb{O} \\ a @ - j \llbracket \varphi' \rrbracket & \text{if } \varphi = \vec{a} \varphi' \text{ and } \vec{a} \text{ is a request action on } a \\ & \text{with } \vec{a}_{(j)} \in \mathbb{R} \\ \varepsilon & \text{if } \varphi = \varepsilon \\ \text{undefined} & \text{otherwise} \end{cases}$$

Now we are ready to establish a first property showing how a sequence of actions labelling a computation of the corresponding communicating system, is mapped into a trace in strong agreement, i.e. a string of matches.

Property 2. Let $S(KS_{\mathcal{A}})$ be the communicating system corresponding to $KS_{\mathcal{A}}$, and let s_0 be its initial configuration. If $s_0 \xrightarrow{f} s$, then

- $f \in \text{Act}^*$ is a sequence of pairs $(\vec{a} @ i j \vec{a} @ i j)$ possibly followed by $\vec{a}' @ i' j'$, for some a, a', i, j, i', j' , and
- there exists a strong agreement φ such that either $f = \llbracket \varphi \rrbracket$ or $f = \llbracket \varphi \rrbracket \vec{a}' @ i' j'$.

Proof. The first item follows from Definition 13, and the second is then immediate by Definition 14. \square

Before establishing our main results, we introduce a notion of well-formedness of contract automata. We require that if an output of a principal i is enabled in two different states, it can be taken in both, giving raise to the same match, regardless of the (projection of these) states on the other principals $j \neq i$ in the product automaton.

In what follows, it is convenient to highlight the principal of a contract automaton that makes an offer or a request. For that, we define $\text{snd}(\vec{a}) \stackrel{\text{def}}{=} i$ when \vec{a} is a match action or an offer action and $\vec{a}_{(i)} \in \mathbb{O}$; similarly, let $\text{rcv}(\vec{a}) \stackrel{\text{def}}{=} j$ when \vec{a} is a match or a request action and $\vec{a}_{(j)} \in \mathbb{R}$. Also, we will omit the target configuration of a transition when immaterial, and simply write, e.g. $\vec{q} \xrightarrow{\vec{a}}$ or $s \xrightarrow{\ell}$.

Definition 15 (Branching Condition). A contract automaton \mathcal{A} has the *branching condition* if and only if the following holds for each \vec{q}_1, \vec{q}_2 reachable in \mathcal{A}

$$\forall \vec{a} \text{ match actions. } (\vec{q}_1 \xrightarrow{\vec{a}} \wedge \text{snd}(\vec{a}) = i \wedge \vec{q}_{1(i)} = \vec{q}_{2(i)}) \text{ implies } \vec{q}_2 \xrightarrow{\vec{a}}$$

Example 8. The product automaton $A \otimes B \otimes C$ of Example 7 enjoys the branching condition. Indeed, consider the match action (\vec{a}, a, \square) and the transition $\vec{q}_0 \xrightarrow{(\vec{a}, a, \square)}$. We have that $\vec{q}_{0(1)} = \vec{q}_{3(1)} = q_{01}$ and there also exists the transition $\vec{q}_3 \xrightarrow{(\vec{a}, a, \square)}$. The same happens for:

$$\begin{array}{lll} \vec{q}_1 \xrightarrow{(\vec{a}, a, \square)}, \vec{q}_4 \xrightarrow{(\vec{a}, a, \square)} & \vec{q}_2 \xrightarrow{(\vec{a}, a, \square)}, \vec{q}_5 \xrightarrow{(\vec{a}, a, \square)} & \vec{q}_6 \xrightarrow{(\vec{a}, a, \square)}, \vec{q}_5 \xrightarrow{(\vec{a}, a, \square)} \\ \vec{q}_7 \xrightarrow{(\vec{a}, a, \square)}, \vec{q}_9 \xrightarrow{(\vec{a}, a, \square)} & \vec{q}_7 \xrightarrow{(\vec{a}, a, \square)}, \vec{q}_8 \xrightarrow{(\vec{a}, a, \square)} & \vec{q}_9 \xrightarrow{(\vec{a}, a, \square)}, \vec{q}_{10} \xrightarrow{(\vec{a}, a, \square)} \\ \vec{q}_8 \xrightarrow{(\vec{a}, a, \square)}, \vec{q}_{10} \xrightarrow{(\vec{a}, a, \square)} & \vec{q}_0 \xrightarrow{(\vec{b}, \square, b)}, \vec{q}_3 \xrightarrow{(\vec{b}, \square, b)} & \vec{q}_0 \xrightarrow{(\square, \vec{c}, c)}, \vec{q}_1 \xrightarrow{(\square, \vec{c}, c)} \end{array}$$

Instead, the most permissive controller of $A \otimes B \otimes C$ does not enjoy the branching condition. Consider again the match action (\vec{a}, a, \square) : while $\vec{q}_{0(1)} = \vec{q}_{3(1)} = q_{01}$ and $\vec{q}_3 \xrightarrow{(\vec{a}, a, \square)}$, no transition labelled by (\vec{a}, a, \square) exits from \vec{q}_0 .

Theorem 1 characterises the relations between a contract automaton \mathcal{A} , its most permissive controller $KS_{\mathcal{A}}$ and the corresponding communicating system $S(KS_{\mathcal{A}})$. It states that $S(KS_{\mathcal{A}})$ is capable of performing all the moves of the controller (item 1), while \mathcal{A} , but possibly not $KS_{\mathcal{A}}$, can perform all the traces of $S(KS_{\mathcal{A}})$ in strong agreement (item 2a). Moreover the runs of $S(KS_{\mathcal{A}})$ leading to a configuration from which no final configuration is reachable correspond to runs in \mathcal{A} that traverse liable transitions (item 2b). The following example illustrates the relations between convergence and deadlock freedom discussed above on the automata in Fig. 5.

Example 9. Consider the contract automata and communicating machines of Fig. 5. A trace of the system $S(KS_{A \otimes B \otimes C})$ is:

$$\begin{aligned} ((q_{01}, q_{02}, q_{03}); (\varepsilon, \varepsilon, \varepsilon)) &\xrightarrow{\vec{a} @ AB} \\ ((q_{11}, q_{02}, q_{03}); (a, \varepsilon, \varepsilon)) &\xrightarrow{a @ AB} \\ ((q_{11}, q_{12}, q_{03}); (\varepsilon, \varepsilon, \varepsilon)) &= s \end{aligned}$$

We have $\llbracket (\bar{a}, a, \square) \rrbracket = \bar{a} @_{AB} a @_{AB}$, and the contract automaton $A \otimes B \otimes C$ is capable of performing the transition $(\bar{q}_0, (\bar{a}, a, \square), \bar{q}_7)$, while this is not true for the controller $KS_{A \otimes B \otimes C}$.

Note that from the configuration s it is not possible to reach a final configuration, since the participant C is prevented from reaching a final state, and thus the transition $(\bar{q}_0, (\bar{a}, a, \square), \bar{q}_7)$ of $A \otimes B \otimes C$ is *liable*.

Moreover s is not a deadlock, indeed it is always possible to perform the loop:

$$\begin{aligned} ((q_{11}, q_{12}, q_{03}); (\varepsilon, \varepsilon, \varepsilon)) &\xrightarrow{\bar{a} @_{AB}} \\ ((q_{11}, q_{12}, q_{03}); (a, \varepsilon, \varepsilon)) &\xrightarrow{a @_{AB}} \\ ((q_{11}, q_{12}, q_{03}); (\varepsilon, \varepsilon, \varepsilon)) & \end{aligned}$$

As a matter of fact $S(KS_{A \otimes B \otimes C})$ is deadlock-free but not convergent.

It is convenient to introduce an *equivalence* between the states of a contract automaton and those of a configuration of a communicating system (assuming them ordered by their indexes).

Definition 16. Let $(\bar{q}; \bar{u})$ be a reachable configuration of a communicating system, and let \bar{q}' be a state of a contract automaton \mathcal{A} . Then we let $\bar{q} \sim \bar{q}'$ iff $|\bar{q}| = |\bar{q}'| = n$ and $\forall i \in 1 \dots n. \bar{q}_{(i)} = \bar{q}'_{(i)}$.

Theorem 1. Let \mathcal{A} be a contract automaton with initial state \bar{q}_0 ; let $KS_{\mathcal{A}}$ be its most permissive controller with initial state \bar{q}_{mpc} ; let $S(KS_{\mathcal{A}})$ be the corresponding communicating system with initial configuration s_0 . Then, given a strong agreement $\varphi \in \mathfrak{Z}$, the following hold:

1. if $\bar{q}_{mpc} \xrightarrow{\varphi} \bar{q}'_{mpc}$, then there exists a configuration s such that $s_0 \xrightarrow{\llbracket \varphi \rrbracket} s = (\bar{q}^s; \bar{u})$ and $\bar{q}'_{mpc} \sim \bar{q}^s$.
2. if $s_0 \xrightarrow{f} s = (\bar{q}^s; \bar{u})$, then
 - (a) if $f = \llbracket \varphi \rrbracket$, then there exists \bar{q}' such that $\bar{q}_0 \xrightarrow{\varphi} \bar{q}'$ and $\bar{q}^s \sim \bar{q}'$
 - (b) if no final configuration is reachable from s , with either $f = \llbracket \varphi \rrbracket$ or $f = \llbracket \varphi \rrbracket \bar{a} @_{ij}$, then the run $\bar{q}_0 \xrightarrow{\hat{\varphi}} \bar{q}'$ has traversed a transition in $TLiable(KS_{\mathcal{A}}/\mathcal{A})$ with either $\hat{\varphi} = \varphi$ or $\hat{\varphi} = \varphi \bar{a}$, where \bar{a} is a match on a and $snd(\bar{a}) = i$, $rcv(\bar{a}) = j$.

Proof. Through the proof assume that \bar{q}' , \bar{q}'_{mpc} and s are such that $\bar{q}_0 \xrightarrow{\varphi} \bar{q}'$, $\bar{q}_{mpc} \xrightarrow{\varphi} \bar{q}'_{mpc}$ and $s_0 \xrightarrow{\llbracket \varphi \rrbracket} s = (\bar{q}^s; \bar{u})$. Also, let \bar{a} be a match with $snd(\bar{a}) = i$ and $rcv(\bar{a}) = j$.

1. By induction on the length of φ .

The base case is when $\bar{q}_{mpc} \xrightarrow{\bar{a}} \bar{q}'_{mpc}$. Let $\bar{q}_{mpc(i)}$ and $\bar{q}_{mpc(j)}$ be the initial states of participants i and j in $S(KS_{\mathcal{A}})$. By Definition 12, we have that

$$(\bar{q}_{mpc(i)}, \bar{a} @_{ij}, \bar{q}'_{mpc(i)}) \quad \text{and} \quad (\bar{q}_{mpc(j)}, a @_{ij}, \bar{q}'_{mpc(j)})$$

are transitions of participants i and j , respectively. We have $s_0 \xrightarrow{\bar{a} @_{ij}} s$ since after the first transition participant j remains in its initial state. Moreover by Definition 12 and Definition 3 we have $\bar{q}'_{mpc(i)} = \bar{q}_{s(i)}$ and $\bar{q}'_{mpc(j)} = \bar{q}_{s(j)}$, all the other components of the states remain in their initial state, and thus $\bar{q}'_{mpc} \sim \bar{q}^s$.

For the inductive case we have $\varphi = \bar{a} \varphi'$, and the run $\bar{q}_{mpc} \xrightarrow{\bar{a}} \bar{q}'_{mpc} \xrightarrow{\varphi'} \bar{q}'_{mpc}$ for some \bar{q}'_{mpc} . The same argument used above guarantees that there exists $s'' = (\bar{q}^{s''}; \bar{u}'')$ and $\bar{q}^{s''} \sim \bar{q}'_{mpc}$, such that $s_0 \xrightarrow{\bar{a} @_{ij}} s'' \xrightarrow{\llbracket \varphi' \rrbracket} s$ and the induction hypothesis suffices.

2. The proof of the statement 2a is by induction on the length of f and the proof of 2b is by contradiction.

- (a) The base case is when $s_0 \xrightarrow{\bar{a} @_{ij}} s$, in other words $s_0 \xrightarrow{\llbracket \bar{a} \rrbracket} s$. Now, in order to obtain $S(KS_{\mathcal{A}})$ through Definition 12 there must exist two automata such that $\bar{q}_{(i)} = \bar{q}_{s(i)}$ and $\bar{q}_{(j)} = \bar{q}_{s(j)}$. Consequently, the product automaton \mathcal{A} has the transition $(\bar{q}_0, \bar{a}, \bar{q})$ where $\forall k \neq i, j. \bar{q}_{0(k)} = \bar{q}_{(k)}$. Thus $\bar{q} \sim \bar{q}^s$.

For the inductive case we have $\varphi = \bar{a} \varphi'$ and the computation $s_0 \xrightarrow{\llbracket \bar{a} \rrbracket} s'' \xrightarrow{\llbracket \varphi' \rrbracket} s$ where $s'' = (\bar{q}^{s''}; \bar{u}'')$. The same argument used above guarantees that there exists a transition $(\bar{q}_0, \bar{a}, \bar{q}'')$ and $\bar{q}'' \sim \bar{q}^{s''}$ and we conclude by applying the induction hypothesis.

- (b) Assume by contradiction that $\bar{q}_0 \xrightarrow{\hat{\varphi}} \bar{q}'$ has traversed no liable transitions. Then by Definition 11 there exists φ' such that $\bar{q}' \xrightarrow{\varphi'} \bar{q}''$ and \bar{q}'' is a final configuration. By Lemma 1 we must have $\bar{q}_{mpc} \xrightarrow{\hat{\varphi} \varphi'} \bar{q}'_{mpc}$ where \bar{q}'_{mpc} is a final configuration. Hence by applying the first item of Theorem 1 we have $s \xrightarrow{f'} s''$ where s'' is a final configuration and $f' = \llbracket \varphi' \rrbracket$ or $f' = \bar{a} @_{ij} \llbracket \varphi' \rrbracket$, obtaining a contradiction. \square

As a side comment to item 2b, we can also prove that if s is a deadlock configuration (and either $f = \llbracket \varphi \rrbracket$ or $f = \llbracket \varphi \rrbracket \bar{a} @ i j$ for some $\varphi \in \mathcal{F}$) and \mathcal{A} enjoys the branching condition, then there exists a run $\vec{q}_0 \xrightarrow{\hat{\varphi}} \vec{q}'$ traversing a transition in $TLiable(KS_{\mathcal{A}}/\mathcal{A})$ with either $\hat{\varphi} = \varphi$ or $\hat{\varphi} = \varphi \bar{a}$, where \bar{a} is a match on a and $snd(\bar{a}) = i$, $rcv(\bar{a}) = j$. Note also that, if a contract automaton \mathcal{A} fires through a liable transition, it can be that $S(KS_{\mathcal{A}})$ never reaches a deadlock configuration, as shown by the [Example 9](#).

We are now ready to state a main result of ours: the controller of a contract automaton has the branching condition if and only if the corresponding communicating system is convergent.

Theorem 2. Let \mathcal{A} be a contract automaton, $KS_{\mathcal{A}}$ be its most permissive controller, and $S(KS_{\mathcal{A}})$ be the corresponding communicating system. Then

$S(KS_{\mathcal{A}})$ is convergent if and only if $KS_{\mathcal{A}}$ satisfies the branching condition.

Proof. Let \vec{q}_0 , \vec{q}_{mpc} , and s_0 be the initial configurations of \mathcal{A} , $KS_{\mathcal{A}}$, and $S(KS_{\mathcal{A}})$, respectively.

(Only if-part) Assume by contradiction that $S(KS_{\mathcal{A}})$ is not convergent, i.e. $s_0 \xrightarrow{f} s = (\vec{q}^s; \vec{u})$ and no final configurations are reachable from s . By [Property 2](#), f is either $f = \llbracket \varphi \rrbracket$ or $f = \llbracket \varphi \rrbracket \bar{a}' @ i' j'$. Hence by applying item 2b of [Theorem 1](#) we have that $\vec{q}_0 \xrightarrow{\hat{\varphi}} \vec{q}'$ has traversed a transition in $TLiable(KS_{\mathcal{A}}/\mathcal{A})$ with either $\hat{\varphi} = \varphi$ or $\hat{\varphi} = \varphi \bar{a}'$. Hence $\hat{\varphi} = \varphi_1 \bar{a} \varphi''$ for some $\varphi_1, \bar{a}, \varphi''$ such that $\vec{q}_0 \xrightarrow{\varphi_1} \vec{q}_1$ is a run of \mathcal{A} and $(\vec{q}_1, \bar{a}, \vec{q}_1') \in TLiable(KS_{\mathcal{A}}/\mathcal{A})$, with $snd(\bar{a}) = i$, $rcv(\bar{a}) = j$ for some $i, j \in \mathcal{P}$.

There exists then a (sub-)computation $s_0 \xrightarrow{\llbracket \varphi_1 \rrbracket} s' = (\vec{q}^{s'}; \vec{u}') \xrightarrow{\bar{a} @ i j} \hat{s}$, and by item 2a of [Theorem 1](#) we have $\vec{q}^{s'} \sim \vec{q}_1$. Since we obtained $S(KS_{\mathcal{A}})$ through [Definition 12](#), the transition $s' \xrightarrow{\bar{a} @ i j} \hat{s}$ has been originated by a transition $(\vec{q}_2, \bar{a}, \vec{q}_3)$ in $KS_{\mathcal{A}}$, for some \vec{q}_2 such that $\vec{q}_{2(i)} = \vec{q}_{1(i)}$. We have that $\vec{q}_1 \neq \vec{q}_2$ because $(\vec{q}_1, \bar{a}, \vec{q}_1') \in TLiable(KS_{\mathcal{A}}/\mathcal{A})$. Hence there must be $\varphi_1 = \varphi_1^1 \varphi_1^2$ such that $\vec{q}_0 \xrightarrow{\varphi_1^1} \vec{q}_2 \xrightarrow{\varphi_1^2} \vec{q}_1$ and the principal (corresponding to participant) i stays idle in each step of φ_1^2 , since $\vec{q}_{2(i)} = \vec{q}_{1(i)}$. The transitions of $KS_{\mathcal{A}}$ include $\vec{q}_2 \xrightarrow{\bar{a}}$, but not $\vec{q}_1 \xrightarrow{\bar{a}}$ and since $\vec{q}_{2(i)} = \vec{q}_{1(i)}$, $KS_{\mathcal{A}}$ violates the branching condition, against our hypothesis.

(If-part) By contradiction assume that the branching condition does not hold in $KS_{\mathcal{A}}$, i.e. in $KS_{\mathcal{A}}$ $\vec{q}_1 \xrightarrow{\bar{a}} \vec{q}'$, $\vec{q}_2 \xrightarrow{\bar{a}}$ where \bar{a} is a match on a with $snd(\bar{a}) = i$, $rcv(\bar{a}) = j$ for some $i, j \in \mathcal{P}$ and $\vec{q}_{1(i)} = \vec{q}_{2(i)}$. Suppose that φ and φ' are such that $\vec{q}_{mpc} \xrightarrow{\varphi} \vec{q}_1$ and $\vec{q}_{mpc} \xrightarrow{\varphi'} \vec{q}_2$.

By the first item of [Theorem 1](#) we have $s_0 \xrightarrow{\llbracket \varphi \rrbracket} s \xrightarrow{\bar{a} @ i j} \hat{s} \xrightarrow{a @ i j} s'' = (\vec{q}^s; \vec{u})$ with $\vec{q}^s \sim \vec{q}'$ and $s_0 \xrightarrow{\llbracket \varphi' \rrbracket} s'$. Moreover, since by hypothesis $\vec{q}_{1(i)} = \vec{q}_{2(i)}$, we also have $s' \xrightarrow{\bar{a} @ i j} \bar{s}$. The computation $\bar{s} \xrightarrow{a @ i j} s_f$, for any φ_2 and s_f final, is not possible, because otherwise by item 2a of [Theorem 1](#) we would have $\vec{q}_2 \xrightarrow{\bar{a} \varphi_2} \vec{q}_f$, with \vec{q}_f final state of the automaton \mathcal{A} , as well of $KS_{\mathcal{A}}$. This would be a contradiction, because $\vec{q}_2 \xrightarrow{\bar{a}}$ by hypothesis. Hence $S(KS_{\mathcal{A}})$ is not convergent, since from \bar{s} it is not possible to reach a final configuration. \square

A consequence of [Theorem 2](#) is that a *strongly safe* contract automaton has the branching condition if and only if its corresponding communicating system is convergent.

Corollary 3. Let \mathcal{A} be a contract automaton, then

\mathcal{A} is strongly safe and satisfies the branching condition if and only if $S(\mathcal{A})$ is convergent.

Proof. The statement follows by applying [Theorem 2](#), because if \mathcal{A} is strongly safe then $\mathcal{A} = KS_{\mathcal{A}}$, hence $KS_{\mathcal{A}}$ has the branching condition. \square


Example 10. Consider the automaton $\mathcal{A} = A \otimes B \otimes C \otimes D$ depicted in [Fig. 6](#). Both principals A and B only offer \bar{a} and then stop, while the others principals B and C only perform the complementary request a .

The contract automaton \mathcal{A} is strongly safe, but it does not enjoy the branching condition. As an example of violation, notice that the state of B in both \vec{q}_1, \vec{q}_3 is the same, i.e. $\vec{q}_{1(2)} = \vec{q}_{3(2)}$. But the match transition $(\square, \bar{a}, \square, a)$ is possibly from state \vec{q}_1 , and not from the state \vec{q}_3 ; also from state \vec{q}_3 we can fire the match transition $(\square, \bar{a}, a, \square)$, which is not available in state \vec{q}_1 .

The translation of \mathcal{A} yields the communicating machines:

$$\begin{aligned} \llbracket KS_{\mathcal{A}} \rrbracket_A &= \bar{a} @ AC + \bar{a} @ AD & \llbracket KS_{\mathcal{A}} \rrbracket_B &= \bar{a} @ BC + \bar{a} @ BD \\ \llbracket KS_{\mathcal{A}} \rrbracket_C &= a @ AC + a @ BC & \llbracket KS_{\mathcal{A}} \rrbracket_D &= a @ AD + a @ BD \end{aligned}$$

A deadlock configuration is generated by the trace $\bar{a} @ AC. a @ AC. \bar{a} @ BC$.

Fig. 6. KS_A .

6. Extensions: agreement and asynchrony

We now extend our proposal along two lines. First we consider the more permissive notion of agreement on contract automata introduced in [3]. Then, we drop the constraints on the number of messages that a buffer can contain, and we consider the fully asynchronous semantics of communicating machines of Definition 6.

6.1. Agreement

The notion of *agreement* of [3] considers as forbidden actions only the non-matching requests while for strong agreement (Definition 7) both non-matching requests and offers are forbidden.

Below, we extend Theorem 2 and establish a correspondence between contract automata enjoying the property of *agreement* [3] and the convergent communicating system with the one buffer semantics. To do that, we introduce a special kind of contract automaton, called *environment*. Then, we extend the branching condition of Definition 15.

We borrow from [3] the definition of agreement and of most permissive controller in this case.

Definition 17 (*Agreement and Controller*). (See [3].) An *agreement* on \mathbb{L} is a finite sequence of match or offer actions. We let \mathfrak{A} to denote the set of all agreements on \mathbb{L} . A contract automaton \mathcal{A} is *safe* if $\mathcal{L}(\mathcal{A}) \subseteq \mathfrak{A}$. We say that \mathcal{A} *admits agreement* when $\mathcal{L}(\mathcal{A}) \cap \mathfrak{A} \neq \emptyset$.

A *controller* of \mathcal{A} is a contract automaton $\mathcal{K}_\mathcal{A}$ such that $\mathcal{L}(\mathcal{K}_\mathcal{A}) \subseteq \mathfrak{A} \cap \mathcal{L}(\mathcal{A})$. The *most permissive controller* of \mathcal{A} is the controller $\mathcal{K}_\mathcal{A}$ such that $\mathcal{L}(\mathcal{K}'_\mathcal{A}) \subseteq \mathcal{L}(\mathcal{K}_\mathcal{A})$ for all $\mathcal{K}'_\mathcal{A}$ controllers of \mathcal{A} .

The construction of the most permissive controller $\mathcal{K}_\mathcal{A}$ is basically the one of Lemma 1 (see for more details [3]), except that only request transitions are removed.

The above notion of agreement does not prevent offer actions to occur in isolation, and the following definition accordingly extends the branching condition of Definition 15.

Definition 18 (*Extended Branching Condition*). A contract automaton \mathcal{A} has the *extended branching condition* if and only if it has the *branching condition* and for each \vec{q}_1, \vec{q}_2 reachable in \mathcal{A} the following holds

$$\forall \vec{a} \text{ offer action } .(\vec{q}_1 \xrightarrow{\vec{a}} \wedge \text{snd}(\vec{a}) = \mathbf{i} \wedge \vec{q}_{1(i)} = \vec{q}_{2(i)}) \text{ implies } \vec{q}_2 \xrightarrow{\vec{a}}$$

To easily handle the offer actions that are now admitted, we introduce a special contract automaton, called *environment*, that “captures” them all. In this way we can re-use the notion of strong agreement. The environment has a single state, both initial and final and a request loop transition for each possible offer.

Definition 19 (*Environment*). The *environment* is the contract automaton

$$\mathbb{E} = \langle \{q_e\}, q_e, \mathbb{L}, \{(q_e, a, q_e) \mid a \in \mathbb{R}\}, \{q_e\} \rangle$$

Note that by composing a contract automaton \mathcal{A} with the environment \mathbb{E} , that is $\mathcal{A} \otimes \mathbb{E}$, all the offer actions of \mathcal{A} are turned into match actions with the environment.

The next theorem shows how the controller of a contract automaton \mathcal{A} and the controller of $\mathcal{A} \otimes \mathbb{E}$ are related by the two branching conditions of Definition 15 and Definition 18.

Theorem 3. *Given the most permissive controller $\mathcal{K}_\mathcal{A}$ for the contract automaton \mathcal{A} , then $\mathcal{K}_\mathcal{A}$ has the extended branching condition if and only if $\mathcal{K}_{\mathcal{A} \otimes \mathbb{E}}$ has the branching condition.*




Fig. 7. From left to right: the contract automaton (with mixed choices) $A \otimes B$ (where $A = \bar{a} + b$ and $B = \bar{b} + a$), and its corresponding communicating machines $\llbracket \mathcal{K}_{A \otimes B} \rrbracket_A$, $\llbracket \mathcal{K}_{A \otimes B} \rrbracket_B$.

Proof.

(if) By hypothesis \mathcal{K}_A has the extended branching condition. By Definition 3 all offers of \mathcal{K}_A are turned into matches with the environment in $\mathcal{K}_{A \otimes E}$. By contradiction assume that $\mathcal{K}_{A \otimes E}$ does not hold the branching condition.

Hence there must be two states \vec{q}_1, \vec{q}_2 and a principal p such that $\vec{q}_1 \xrightarrow{\bar{a}}, \vec{q}_2 \xrightarrow{\bar{a}}$ with $\vec{q}_1(p) = \vec{q}_2(p)$ where \bar{a} is a match with $\text{snd}(\bar{a}) = p$.

Moreover it must be that $\text{rcv}(\bar{a}) = E$ (by abuse of notation E is the principal corresponding to the environment); otherwise the match transition would also be in \mathcal{K}_A , obtaining a contradiction. By Definition 3 the match \bar{a} in $\mathcal{K}_{A \otimes E}$ is turned into an offer action \bar{a}' in \mathcal{K}_A , and we have that \mathcal{K}_A does not hold the extended branching condition since $\vec{q}_1 \xrightarrow{\bar{a}'}, \vec{q}_2 \xrightarrow{\bar{a}'}$.

(only if) By hypothesis we have that $\mathcal{K}_{A \otimes E}$ has the branching condition. By Definition 3 we know that in the contract automaton $\mathcal{K}_{A \otimes E}$ all the matches involving the environment are coupled with offers of \mathcal{K}_A .

By contradiction assume that \mathcal{K}_A does not hold the extended branching condition. Hence there must be two states \vec{q}_1, \vec{q}_2 and a principal p such that $\vec{q}_1 \xrightarrow{\bar{a}}, \vec{q}_2 \xrightarrow{\bar{a}}$ with $\vec{q}_1(p) = \vec{q}_2(p)$ and $\text{snd}(\bar{a}) = p$.

We distinguish two cases:

- \bar{a} is a match action: then by Definition 3 the match is present also in $\mathcal{K}_{A \otimes E}$, obtaining a contradiction since $\mathcal{K}_{A \otimes E}$ has the branching condition.
- \bar{a} is an offer action: then by Definition 3 \bar{a} is turned into a match \bar{a}' with E in $\mathcal{K}_{A \otimes E}$. Since $\vec{q}_1 \xrightarrow{\bar{a}'}, \vec{q}_2 \xrightarrow{\bar{a}'}$ we have that $\mathcal{K}_{A \otimes E}$ does not hold the branching condition, obtaining a contradiction. \square

Finally we relate a controller \mathcal{K}_A of a contract automaton \mathcal{A} with the corresponding system $S(\mathcal{K}_{A \otimes E})$, obtained from the controller composed with the environment. Indeed as a direct consequence of Theorem 2 and Theorem 3 we have that \mathcal{K}_A has the extended branching condition if and only if $S(\mathcal{K}_{A \otimes E})$ is convergent.

Corollary 4. Given the most permissive controller \mathcal{K}_A for the contract automaton \mathcal{A} , then \mathcal{K}_A has the extended branching condition iff $S(\mathcal{K}_{A \otimes E})$ is convergent.

Proof. (if) By hypothesis \mathcal{K}_A has the extended branching condition, by Theorem 3 we have that $\mathcal{K}_{A \otimes E}$ has the branching condition. Finally by Theorem 2 we have that $S(\mathcal{K}_{A \otimes E})$ is convergent.

(only if) By hypothesis $S(\mathcal{K}_{A \otimes E})$ is convergent, by applying Theorem 2 we have that $\mathcal{K}_{A \otimes E}$ has the branching condition. Finally by Theorem 3 we have that \mathcal{K}_A has the extended branching condition. \square

6.2. Asynchronous semantics of communicating systems

We now discuss the relations between contract automata and communicating systems, with the semantics of Definition 6, i.e. when the buffers of the communicating machines can contain more than one message. From now onward, the notions of *convergence* and *deadlock* introduced in Definition 13 are considered using the semantics of Definition 6. The following example shows the difficulties with the unrestricted semantics.

Example 11. Consider the contract automaton $A \otimes B$ and the corresponding communicating machines of Fig. 7. This contract automaton is strongly safe and has the branching condition. However the translated system is *not* convergent. Indeed a possible deadlock in $S(\mathcal{K}_{A \otimes B})$ occurs if A performs the action $\bar{a} @ AB$ and then B performs the action $\bar{b} @ BA$. This is because participant B can ignore the message received by the participant A and follow the other branch of the contract automaton. These behaviours are not permitted by the 1-buffer semantics.

In order to guarantee that a system of communicating machines corresponding to a contract automaton is convergent, we constrain the contract automata as follows. Intuitively, we will discard those contract automata that have a request and an offer transition of a principal outgoing from the same node, like the automata A and B of Example 11. More precisely:

Definition 20. Let \mathcal{A} be a contract automaton, then \mathcal{A} has a *mixed choice* if and only if there exists a reachable state \vec{q} with two outgoing transitions $\vec{q} \xrightarrow{\vec{a}_1}, \vec{q} \xrightarrow{\vec{a}_2}$ such that $\text{snd}(\vec{a}_1) = \text{rcv}(\vec{a}_2)$.

In the following we will prove that if the strong controller of a contract automaton \mathcal{A} has the branching condition and no mixed choices then the corresponding system is convergent.

It is convenient to define how to project a trace of a communicating system into its offer actions, and to prove an auxiliary property.

Definition 21. Given $f \in \text{Act}^*$, its projection on its offers is defined as follows:

$$f|_{\mathbb{O}} \stackrel{\text{def}}{=} \begin{cases} \vec{a}@i:j(f'|_{\mathbb{O}}) & \text{if } f = \vec{a}@i:j f' \text{ for some } i, j \in \mathcal{P}, \vec{a} \in \mathbb{O} \\ f'|_{\mathbb{O}} & \text{if } f = a@i:j f' \text{ for some } i, j \in \mathcal{P}, a \in \mathbb{R} \\ \varepsilon & \text{if } f = \varepsilon \end{cases}$$

By abuse of notation, we say that a communicating machine has no mixed choices if it is never the case that both offer and request transitions leave one of its states.

Property 5. Let $KS_{\mathcal{A}}$ be the most permissive controller of \mathcal{A} and $S(KS_{\mathcal{A}})$ be the corresponding communicating system.

If $KS_{\mathcal{A}}$ has the branching condition and no mixed choices then all the communicating machines of the participants in $S(KS_{\mathcal{A}})$ have no mixed choices.

Proof. By contradiction assume that there is a participant p with two transitions $(q_1, \vec{a}@p:j, q_2), (q_1, b@i:p, q_3)$.

By Definition 12 there must be two transitions $(\vec{q}, \vec{a}, \vec{q}_1), (\vec{q}, \vec{b}, \vec{q}_3)$ where $\vec{q}_{(p)} = \vec{q}_{2(p)}$, and \vec{a}, \vec{b} are match actions on a and b , respectively, with $\text{snd}(\vec{a}) = p, \text{rcv}(\vec{a}) = j, \text{snd}(\vec{b}) = i, \text{rcv}(\vec{b}) = p$.

There are the following two cases:

- $\vec{q} = \vec{q}_2$ we obtain a contradiction since we have a mixed choice;
- $\vec{q} \neq \vec{q}_2$ then since $\vec{q}_{(p)} = \vec{q}_{2(p)}$ by the branching condition there are also two transitions $(\vec{q}_2, \vec{a}, \vec{q}_4), (\vec{q}, \vec{b}, \vec{q}_5)$, hence both \vec{q}, \vec{q}_2 are mixed choices, obtaining a contradiction. \square

The following theorem relates the traces of a communicating system with the ones of the strong controller $KS_{\mathcal{A}}$ it comes from, provided that $KS_{\mathcal{A}}$ has the branching condition and no mixed choices. Under the above conditions a trace of a communicating system only differs from the corresponding trace of the $KS_{\mathcal{A}}$ in the order in which the request actions are fired in the system. Indeed by considering only offer actions the two traces are *equal*.

Theorem 4. Let $KS_{\mathcal{A}}$ be the most permissive controller of \mathcal{A} , with the branching condition and no mixed choices, and let Φ be the set of its non-empty traces; let $S(KS_{\mathcal{A}})$ be the communicating system obtained from $KS_{\mathcal{A}}$ and let F be the set of its non-empty traces. Then

$$\forall f \in F \text{ there exists } \varphi \in \Phi \text{ such that } f|_{\mathbb{O}} = \llbracket \varphi \rrbracket|_{\mathbb{O}}.$$

Proof. Let \vec{q}_{mpc} and s be the initial configurations of $KS_{\mathcal{A}}$ and $S(KS_{\mathcal{A}})$, respectively. Assume by contradiction that there exists a f with $s \xrightarrow{f}$ such that for all φ with $\vec{q}_{mpc} \xrightarrow{\varphi}$ we have $f|_{\mathbb{O}} \neq \llbracket \varphi \rrbracket|_{\mathbb{O}}$.

Since $f \neq \varepsilon$, by Definition 6 it must be that $f = \vec{a}_1@i_1:j_1 f'_1$ for some $i_1, j_1, \vec{a}_1, f'_1$, and by Definition 12 there must be a transition $(\vec{q}_{mpc}, \vec{a}', \vec{q}'_{mpc})$ in $KS_{\mathcal{A}}$, and therefore there is a $\varphi = \vec{a}'\varphi'_1$, for some \vec{q}'_{mpc} and φ'_1 , where \vec{a}' is a match on a_1 with $\text{snd}(\vec{a}') = i_1, \text{rcv}(\vec{a}') = j_1$. Moreover by hypothesis it must be that $f'_1|_{\mathbb{O}} \neq \llbracket \varphi'_1 \rrbracket|_{\mathbb{O}}$.

Now split f as $f_1 f_2$, and select a trace $\varphi = \varphi_1 \varphi_2$ such that f_1 is the longest prefix of f with $\llbracket \varphi_1 \rrbracket|_{\mathbb{O}} = f_1|_{\mathbb{O}} (\neq \varepsilon$ because of the above). Then we have $f_2 = \vec{a}@i:j f_3$ for some i, j, \vec{a} , and $\vec{q}_{mpc} \xrightarrow{\varphi_1} \vec{q}_{1mpc} \xrightarrow{\vec{a}} \vec{q}_r$ where \vec{a} is a match with $\text{snd}(\vec{a}) = i, \text{rcv}(\vec{a}) = j$. Then, there must exist a transition $(\vec{q}_{2mpc}, \vec{a}, \vec{q}'_{2mpc})$, for some $\vec{q}_{2mpc}, \vec{q}'_{2mpc}$. Assuming $s \xrightarrow{f_1} (\vec{q}_r, \vec{u})$, we have then that $\vec{q}_{r(i)} = \vec{q}_{2mpc(i)}$.

Since the branching condition holds, it turns out that $\vec{q}_{1mpc(i)} \neq \vec{q}_{2mpc(i)}$. If the principal (corresponding to participant) i stays idle in φ_1 then it must be $\vec{q}_{mpc(i)} = \vec{q}_{1mpc(i)} = \vec{q}_{2mpc(i)}$ obtaining a contradiction. Thus, the principal i has performed some steps in φ_1 , more precisely some requests, otherwise we would obtain $\vec{q}_{1mpc(i)} = \vec{q}_{2mpc(i)}$, because $\llbracket \varphi_1 \rrbracket|_{\mathbb{O}} = f_1|_{\mathbb{O}}$.

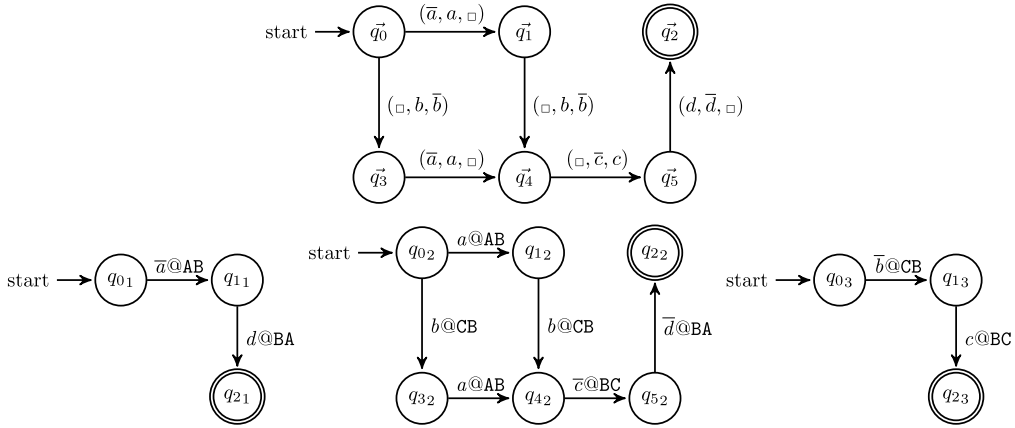


Fig. 8. Top: $KS_{A \otimes B \otimes C}$. Bottom: $\llbracket KS_{A \otimes B \otimes C} \rrbracket_A$, $\llbracket KS_{A \otimes B \otimes C} \rrbracket_B$, $\llbracket KS_{A \otimes B \otimes C} \rrbracket_C$.

We now introduce an auxiliary notion of projection on the offers of a participant p :

$$\hat{f}|_{\mathcal{O}, p} \stackrel{\text{def}}{=} \begin{cases} \bar{a}@p j \hat{f}'|_{\mathcal{O}, p} & \text{if } \hat{f} = \bar{a}@p j \hat{f}' \text{ for some } j \in \mathcal{P}, \bar{a} \in \mathbb{O} \\ \hat{f}'|_{\mathcal{O}, p} & \text{if } \hat{f} = a@i j \hat{f}' \text{ for some } i, j \in \mathcal{P}, a \in \mathbb{R} \\ \hat{f}'|_{\mathcal{O}, p} & \text{if } \hat{f} = \bar{a}@i j \hat{f}' \text{ for some } i, j \in \mathcal{P}, i \neq p, \bar{a} \in \mathbb{O} \\ \varepsilon & \text{if } \hat{f} = \varepsilon \end{cases}$$

We consider the offers of i , if any.

- $\llbracket \varphi_1 \rrbracket|_{\mathcal{O}, i} = f_1|_{\mathcal{O}, i} = \varepsilon$: if all the requests $a@i j$ occurring in (matches of) φ_1 also occur in f_1 , then we obtain $\bar{q}_{1_{mpc(i)}} = \bar{q}_{2_{mpc(i)}}$, and there is a violation of the branching condition. Hence not all the requests of i in φ_1 occur in f_1 , i.e. $\bar{q}_{mpc} \xrightarrow{\varphi_1^1} \bar{q}_{2_{mpc}} \xrightarrow{\varphi_1^2} \bar{q}_{1_{mpc}}$ with $\varphi_1 = \varphi_1^1 \varphi_1^2$ and assume that the missing requests occur in φ_1^2 . Then further split φ_1^2 as $\varphi_1^3 \bar{a}_2 \varphi_1^4$, where $rcv(\bar{a}_2) = i$, so that $\bar{q}_{2_{mpc}} \xrightarrow{\varphi_1^3} \bar{q}_{4_{mpc}} \xrightarrow{\bar{a}_2 \varphi_1^4} \bar{q}_{1_{mpc}}$ and that principal i stays idle in φ_1^1 . Then we have that $\bar{q}_{2_{mpc(i)}} = \bar{q}_{4_{mpc(i)}}$ and by the branching condition we have the transition $(\bar{q}_{4_{mpc}}, \bar{a}, \bar{q}_{4'_{mpc}})$ showing a mixed choice, because the principal i can perform both an offer and a request: contradiction.
- $\llbracket \varphi_1 \rrbracket|_{\mathcal{O}, i} = f_1|_{\mathcal{O}, i} \neq \varepsilon$: assume that the last offer of i in $\llbracket \varphi_1 \rrbracket$ is $\bar{a}_5@i j_2$, i.e. $\llbracket \varphi_1 \rrbracket|_{\mathcal{O}, i} = \llbracket \varphi_1^1 \rrbracket|_{\mathcal{O}, i} \bar{a}_5@i j_2$, with $\llbracket \varphi_1 \rrbracket = \llbracket \varphi_1^1 \rrbracket \bar{a}_5@i j_2 \llbracket \varphi' \rrbracket$ and $\llbracket \varphi' \rrbracket|_{\mathcal{O}, i} = \varepsilon$. Also, let $\bar{q}_{mpc} \xrightarrow{\varphi_1^1 \bar{a}_5} \bar{q}_{5_{mpc}}$ where \bar{a}_5 is a match on a_5 with $snd(\bar{a}_5) = i$, $rcv(\bar{a}_5) = j_2$. Now, we show that any requests of i occurring in $\llbracket \varphi_1^1 \rrbracket$, also occur in f_1 before the offer $\bar{a}_5@i j_2$. Indeed, by [Property 5](#) we have that the communicating machine of the participant i has no mixed choices (so it cannot choose between offering or requesting) and is therefore forced to read from its buffer the received offers before firing $\bar{a}_5@i j_2$, because these were coupled (with corresponding requests of i) in matches occurring in $\llbracket \varphi_1^1 \rrbracket$. We conclude the proof by applying the argument of the previous case with $\bar{q}_{5_{mpc}}$ as the initial state, given that $\llbracket \varphi' \rrbracket|_{\mathcal{O}, i} = \varepsilon$. \square

Note in passing that for simulating a step in non-empty trace φ of a most permissive controller, the corresponding communicating system can pass through several possible configurations, in order to execute a sequence of actions f such that $f|_{\mathcal{O}} = \llbracket \varphi \rrbracket|_{\mathcal{O}}$.

As a matter of fact, a participant can fire many requests, not registered in $\llbracket \varphi \rrbracket|_{\mathcal{O}}$, after its last offer registered therein. Indeed the trace φ is formed by match actions, hence all the requests are fired, while in the trace f there could be some participant which has fired all the offers but not yet all its requests.

Example 12. Consider the automata depicted in [Fig. 8](#). The controller $KS_{A \otimes B \otimes C}$ has the branching condition and no mixed choices. Three possible traces of $S(KS_{A \otimes B \otimes C})$ are:

$$f = \bar{a}@AB \bar{b}@CB a@AB b@CB \bar{c}@BC \bar{d}@BA \quad f_1 = f \ c@BC \quad f_2 = f_1 \ d@BA$$

Consider the trace of $KS_{A \otimes B \otimes C}$:

$$\varphi = (\bar{a}, a, \square)(\square, b, \bar{b})(\square, \bar{c}, c)(d, \bar{d}, \square)$$




Fig. 9. From left to right: the contract automaton (with mixed choices) $A \otimes B$ (where $A = \bar{a}.b + b.\bar{a}$ and $B = \bar{b}.a + a.\bar{b}$), and its corresponding communicating machines $\llbracket \mathcal{K}_{A \otimes B} \rrbracket_A$, $\llbracket \mathcal{K}_{A \otimes B} \rrbracket_B$.

We have $f \downarrow_0 = f_1 \downarrow_0 = f_2 \downarrow_0 = \llbracket \varphi \rrbracket \downarrow_0$. Note that in order to fire the offer $\bar{c} @ BC$ the participant C needs to fire all the previous requests. Moreover in the trace f_2 each participant has fired all its requests, indeed we have: $\vec{q}_{mpc} \xrightarrow{\varphi} \vec{q}'_{mpc}$, $s \xrightarrow{f_2} (\vec{q}_r; \vec{u})$ and $\vec{q}'_{mpc} \sim \vec{q}_r$, where \vec{q}_{mpc} and s are the starting configurations of $KS_{A \otimes B \otimes C}$ and $S(KS_{A \otimes B \otimes C})$.

The main result of this sub-section follows. If the strong controller of a contract automaton \mathcal{A} has the branching condition and no mixed choices, then its corresponding system is convergent with unrestricted semantics.

Theorem 5. Let $S(KS_{\mathcal{A}})$ be the communicating system corresponding to the most permissive controller $KS_{\mathcal{A}}$. Then

$KS_{\mathcal{A}}$ has the branching condition and no mixed choices implies $S(KS_{\mathcal{A}})$ is convergent.

Proof. Let s be the initial configuration of $S(KS_{\mathcal{A}})$ and \vec{q}_{mpc} be the initial state of $KS_{\mathcal{A}}$. By contradiction assume that $KS_{\mathcal{A}}$ has the branching condition and no mixed choices but $S(KS_{\mathcal{A}})$ is not convergent, i.e. there exists a non-empty run f such that $s \xrightarrow{f} (\vec{q}_r; \vec{u})$ and no final configurations are reachable from $(\vec{q}_r; \vec{u})$. By Theorem 4 there exists a non-empty trace of the controller φ such that $\vec{q}_{mpc} \xrightarrow{\varphi} \vec{q}'_{mpc}$ and $f \downarrow_0 = \llbracket \varphi \rrbracket \downarrow_0$. Since all the offers are matched in φ , by Definitions 6 and 12 there exists a trace where all the request actions occur, and let it be our f . Since both $KS_{\mathcal{A}}$ and $S(KS_{\mathcal{A}})$ are deterministic, and since f drives each communicating machine to execute all and only the actions occurring in the matches of φ , we have that $\vec{q}_r \sim \vec{q}'_{mpc}$.

Finally since $KS_{\mathcal{A}}$ is a controller, there must be a trace φ' and a final state \vec{q}_f such that $\vec{q}'_{mpc} \xrightarrow{\varphi'} \vec{q}_f$. By applying Theorem 1.1 we obtain $(\vec{q}_r; \vec{u}) \xrightarrow{\llbracket \varphi' \rrbracket} s_f$ where s_f is a final configuration, because the transition relation of the 1-buffer semantics is included in that of the semantics of Definition 6: contradiction. \square

Note that the above theorem does not fully generalise Corollary 3, as shown by the following example.

Example 13. Consider the contract automaton of Fig. 9 that is strongly safe (so it is also its most permissive controller) and has the branching condition. Its corresponding system consists of the communicating machines in Fig. 9 and it is convergent. However, a mixed choice is possible from $\vec{q}_0 = (q_{01}, q_{02})$.

7. An example

We show our proposal at work on the *two buyers protocol* (2BP for short) presented in [31]. There are two buyers B_1 and B_2 that collaborate in purchasing an item from a seller S . Buyer B_1 starts the protocol by asking S the price of the desired item (*price*); the seller S replies with the quote for the requested item by sending the quotation message *quote* to both buyers. Once received its quote, buyer B_1 sends to B_2 its contribution for purchasing the item (*contrib*).

Buyer B_2 waits for the quote from S and the contribution from B_1 . Then, it decides whether to terminate by issuing the *noop* message to S , or to proceed by sending an acknowledgement to S .

Upon receiving the acknowledgement, the seller sends the item to B_2 (*delivery*), while if it receives *noop* it terminates with no further action.

The contract automata B_1 , B_2 and S admit strong agreement, and they are shown in Fig. 10, together with their most permissive controller.

We now analyse the translation of the orchestration into a choreography of communicating finite state machines. The most permissive controller does not enjoy the branching condition because, for example:

- in \vec{q}_2 and \vec{q}_4 , buyer B_1 is in state q_{B12} ,
- $\vec{q}_4 \xrightarrow{(\text{contrib}, \text{contrib}, \square)} \dots$, but
- there is no transition from \vec{q}_2 such that $\vec{q}_2 \xrightarrow{(\text{contrib}, \text{contrib}, \square)} \dots$,

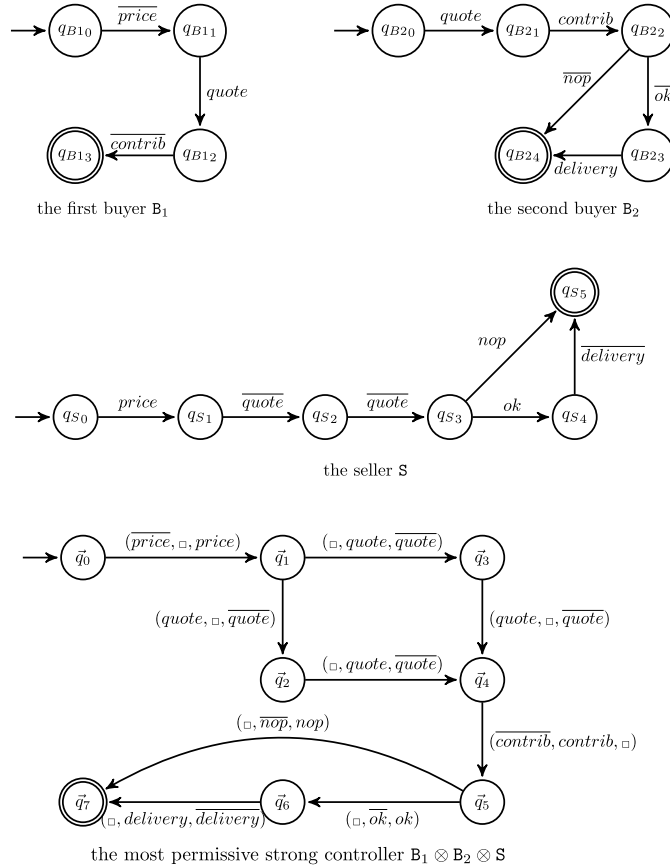


Fig. 10. The contract automata for 2BP and their most permissive controller.

as required by Definition 15. This happens because B_1 could send the contribution to B_2 before this has received *quote* from S , so blocking the system. An easy recover would be forcing the seller to first send the quotation to buyer B_2 , so reducing the nondeterminism. A convergent choreography of communicating machines, both for the synchronous and the asynchronous case is then derivable as specified above.

The most permissive controller has no states with mixed choices. However, this case may arise if B_2 could additionally withdraw before accepting the contribution from B_1 , because the quotation is too high. It is out of the scope of this paper discussing the ways to recover from this situation, and we only note that our proposal clearly detects why and where it shows up, making the corresponding choreography not convergent.

8. Concluding remarks

We have established a formal correspondence between contract automata, an orchestration model, and communicating systems, i.e. sets of communicating machines, that is a model of choreography. More precisely, we introduced the notion of strong agreement on contract automata, stating that all the traces accepted by the automata are made by requests matched by corresponding offers; on the second model we considered convergence, i.e. successful termination of computations. We proved that strong agreement corresponds to convergence (cf. Theorem 2) when communicating systems are endowed with the 1-buffer semantics, according to which the execution of the machines is basically synchronous. We note in passing that this has some advantages since communicating systems with the 1-buffer semantics are computationally more tractable than in the general case [30]. We then generalised the above result by adopting a more relaxed notion of agreement that admits computations where offer actions can go unmatched (cf. Corollary 4). We also proved a slightly weaker result for communicating systems with the unrestricted semantics (cf. Theorem 5). Currently, we are developing a verification toolkit based on the results presented here and in [3]; a prototype is available at <https://github.com/davidebasile/workspace>.

The constructions to obtain contract automata from principals (e.g., products of contract automata, controllers, controlled systems) could be complemented with the recent results on the synthesis of choreographies for communicating finite state machines proposed by [11]. We conjecture that the conditions guaranteeing well-behaviour identified here imply the *generalised multiparty compatibility* condition identified in [11]. For a given communicating system satisfying this property, a choreography always exists and can be synthesised effectively.

In [32] progress is attained under assumptions allowing a process P to expose unmatched actions as long as the closed system made of P and a *catalyser* process (derived from P) results to be lock-free. We can give an interesting analogy with this work: our environments resemble their *catalysers* and the extended branching condition (cf. Definition 13) is similar to their relaxed notion of safety. As far as we can see, our notion of convergence is stronger than the above notion of safety, in that we also guarantee that a final state is reached, possibly infinitely often. Additionally, only dyadic session types are dealt with in [32], while here we can cope with multiparty scenarios.

A possible application of the results in [11] is to use the projections obtained from the communicating system corresponding to the controller. Through them, we can identify the principals that could be liable. Indeed, each principal univocally corresponds to one of those projections. Hence, this would allow to flag the components that may lead to communication mismatches so to refine them and guarantee that the refined principals execute without the intervention of an orchestrator. In general, the suggested approach would give a more efficient and more distributed execution. This is because one can remove the overhead due to the communication with the orchestrator, and avoid the centralisation point of the orchestrator, respectively. Additionally, one could find that only few principals, and more importantly which of them, spoil the conditions for achieving choreographed executions. Hence, by modifying/replacing those principals with the machines obtained by projecting the synthesised choreography it would be possible to retrieve a choreographed executions of the contract automata.

References

- [1] C. Peltz, Web services orchestration and choreography, *IEEE Comput.* 36 (10) (2003) 46–52.
- [2] N. Kavantzaz, D. Burdett, G. Ritzinger, T. Fletcher, Y. Lafon, C. Barreto, Web services choreography description language version 1.0, World Wide Web Consortium, Candidate Recommendation CR-ws-cdl-10-20051109, November 2005.
- [3] D. Basile, P. Degano, G.L. Ferrari, Automata for analysing service contracts, in: M. Maffei, E. Tuosto (Eds.), Trustworthy Global Computing – 9th International Symposium, Revised Selected Papers, TGC 2014, Rome, Italy, in: Lecture Notes in Computer Science, vol. 8902, Springer, 2014, pp. 34–50.
- [4] D. Brand, P. Zafiropulo, On communicating finite-state machines, *J. ACM* 30 (2) (1983) 323–342.
- [5] M. Bartoletti, E. Tuosto, R. Zunino, On the realizability of contracts in dishonest systems, in: M. Sirjani (Ed.), Coordination Models and Languages – 14th International Conference, COORDINATION 2012, Proceedings, in: Lecture Notes in Computer Science, vol. 7274, Springer, 2012, pp. 245–260.
- [6] M. Bartoletti, A. Scalas, E. Tuosto, R. Zunino, Honesty by typing, in: D. Beyer, M. Boreale (Eds.), Formal Techniques for Distributed Systems – Joint IFIP WG 6.1 International Conference, FMOODS/FORTE 2013, Proceedings, in: Lecture Notes in Computer Science, vol. 7892, Springer, 2013, pp. 305–320.
- [7] P. Deniérou, N. Yoshida, Multiparty session types meet communicating automata, in: H. Seidl (Ed.), Programming Languages and Systems – 21st European Symposium on Programming, ESOP 2012, Proceedings, in: Lecture Notes in Computer Science, vol. 7211, Springer, 2012, pp. 194–213.
- [8] C.G. Cassandras, S. LaFortune, Introduction to Discrete Event Systems, Springer, Secaucus, NJ, USA, 2006.
- [9] D. Basile, P. Degano, G.L. Ferrari, E. Tuosto, From orchestration to choreography through contract automata, in: I. Lanese, A. Lluch-Lafuente, A. Sokolova, H.T. Vieira (Eds.), Proceedings 7th Interaction and Concurrency Experience, ICE 2014, in: EPTCS, vol. 166, 2014, pp. 67–85.
- [10] I. Lanese, C. Guidi, F. Montesi, G. Zavattaro, Bridging the gap between interaction- and process-oriented choreographies, in: Software Engineering and Formal Methods, SEFM 2008, 2008, pp. 323–332.
- [11] J. Lange, E. Tuosto, N. Yoshida, From communicating machines to graphical choreographies, in: S.K. Rajamani, D. Walker (Eds.), Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, ACM, 2015, pp. 221–232.
- [12] J. Lange, A. Scalas, Choreography synthesis as contract agreement, in: Proceedings 6th Interaction and Concurrency Experience, ICE 2013, 2013, pp. 52–67.
- [13] R. Dijkman, M. Dumas, Service-oriented design: a multi-viewpoint approach, *Int. J. Coop. Inf. Syst.* 13 (4) (2004) 337–368.
- [14] G. Alonso, F. Casati, H. Kuno, V. Machiraju, Web Services – Concepts, Architectures and Applications. Data-Centric Systems and Applications, Springer, 2004.
- [15] M. Papazoglou, Web Services and SOA: Principles and Technology, Pearson-Prentice Hall, 2012.
- [16] J. Misra, Computation orchestration, in: M. Broy, J. Grünbauer, D. Harel, T. Hoare (Eds.), Engineering Theories of Software Intensive Systems, in: NATO Science Series, vol. 195, Springer, 2005, pp. 285–330.
- [17] Z. Qiu, X. Zhao, C. Cai, H. Yang, Towards the theoretical foundation of choreography, in: C.L. Williamson, M.E. Zurko, P.F. Patel-Schneider, P.J. Shenoy (Eds.), Proceedings of the 16th International Conference on World Wide Web, WWW 2007, ACM, 2007, pp. 973–982.
- [18] H. Yang, X. Zhao, C. Cai, Z. Qiu, Exploring the connection of choreography and orchestration with exception handling and finalization/compensation, in: J. Derrick, J. Vain (Eds.), Formal Techniques for Networked and Distributed Systems – FORTE 2007, 27th IFIP WG 6.1 International Conference, Proceedings, in: LNCS, vol. 4574, Springer, 2007, pp. 81–96.
- [19] S.R. Talbot, Orchestration and choreography: standards, tools and technologies for distributed workflows, http://www.nettab.org/2005/docs/NETTAB2005_Ross-TalbotOral.pdf.
- [20] N. Busi, R. Gorrieri, C. Guidi, R. Lucchi, G. Zavattaro, Choreography and orchestration conformance for system design, in: P. Ciancarini, H. Wiklicky (Eds.), Coordination Models and Languages, 8th International Conference, COORDINATION 2006, Proceedings, in: Lecture Notes in Computer Science, vol. 4038, Springer, 2006, pp. 63–81.
- [21] I. Lanese, F. Montesi, G. Zavattaro, Amending choreographies, in: A. Ravara, J. Silva (Eds.), Proceedings 9th International Workshop on Automated Specification and Verification of Web Systems, WWV 2013, in: EPTCS, vol. 123, 2013, pp. 34–48.
- [22] S. Basu, T. Bultan, M. Quederni, Deciding choreography realizability, in: J. Field, M. Hicks (Eds.), Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, ACM, 2012, pp. 191–202.
- [23] S. Hallé, T. Bultan, Realizability analysis for message-based interactions using shared-state projections, in: G. Roman, K.J. Sullivan (Eds.), Proceedings of the 18th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2010, ACM, 2010, pp. 27–36.
- [24] C. Laneve, L. Padovani, An algebraic theory for web service contracts, *Form. Asp. Comput.* (2015) 1–28.
- [25] S. Basu, T. Bultan, Choreography conformance via synchronizability, in: S. Srinivasan, K. Ramamritham, A. Kumar, M.P. Ravindra, E. Bertino, R. Kumar (Eds.), Proceedings of the 20th International Conference on World Wide Web, WWW 2011, ACM, 2011, pp. 795–804.
- [26] M. Bravetti, G. Zavattaro, Contract-based discovery and composition of web services, in: M. Bernardo, L. Padovani, G. Zavattaro (Eds.), Formal Methods for Web Services, 9th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, 2009, in: Lecture Notes in Computer Science, vol. 5569, Springer, 2009, pp. 261–295.
- [27] M. Bravetti, G. Zavattaro, Contract compliance and choreography conformance in the presence of message queues, in: R. Bruni, K. Wolf (Eds.), Web Services and Formal Methods, 5th International Workshop, WS-FM 2008, Revised Selected Papers, in: Lecture Notes in Computer Science, vol. 5387, Springer, 2008, pp. 37–54.

- [28] M. Autili, D.D. Ruscio, A.D. Salle, P. Inverardi, M. Tivoli, A model-based synthesis process for choreography realizability enforcement, in: V. Cortellessa, D. Varró (Eds.), *Fundamental Approaches to Software Engineering – 16th International Conference, FASE 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, 2013*. Proceedings, in: *Lecture Notes in Computer Science*, vol. 7793, Springer, 2013, pp. 37–52.
- [29] L. de Alfaro, T.A. Henzinger, Interface automata, in: *Proceedings of the 8th European Software Engineering Conference Held Jointly with 9th ACM SIGSOFT International Symposium on Foundations of Software Engineering 2001*, ACM, 2001, pp. 109–120.
- [30] G. Cécé, A. Finkel, Verification of programs with half-duplex communication, *Inf. Comput.* 202 (2) (2005) 166–190.
- [31] K. Honda, N. Yoshida, M. Carbone, Multiparty asynchronous session types, in: G.C. Necula, P. Wadler (Eds.), *POPL*, ACM, 2008, pp. 273–284.
- [32] M. Carbone, O. Dardha, F. Montesi, Progress as compositional lock-freedom, in: E. Kühn, R. Pugliese (Eds.), *COORDINATION 2014*, in: *Lecture Notes in Computer Science*, vol. 8459, Springer, 2014, pp. 49–64.