Towards Dynamic Updates in Service Composition

Mario Bravetti
University of Bologna, Italy / INRIA, France
mario.bravetti@unibo.it

We survey our results about verification of adaptable processes. We present adaptable processes as a way of overcoming the limitations that process calculi have for describing patterns of dynamic process evolution. Such patterns rely on direct ways of controlling the behavior and location of running processes, and so they are at the heart of the adaptation capabilities present in many modern concurrent systems. Adaptable processes have named scopes and are sensible to actions of dynamic update at runtime; this allows to express dynamic and static topologies of adaptable processes as well as different evolvability patterns for concurrent processes. We introduce a core calculus of adaptable processes and consider verification problems for them: first based on specific properties related to error occurrence, that we call bounded and eventual adaptation, and then by considering a simple yet expressive temporal logic over adaptable processes. We provide (un)decidability results of such verification problems over adaptable processes considering the spectrum of topologies/evolvability patterns introduced. We then consider distributed adaptability, where a process can update part of a protocol by performing dynamic distributed updates over a set of protocol participants. Dynamic updates in this context are presented as an extension of our work on choreographies and behavioural contracts in multiparty interactions. We show how update mechanisms considered for adaptable processes can be used to extend the theory of choreography and orchestration/contracts, allowing them to be modified at run-time by internal (self-adaptation) or external intervention.

1 Introduction

We survey our work about verification of adaptable processes [2, 3, 1] presenting theories that we previously introduced in different contexts and connecting them for the first time. We start from adaptation mechanisms where updates take place on individual processes only and we classify them according to update patterns and adaptable process topologies. Within such classification we present separation results based un(decidability) of verification of adaptation properties/sets of temporal logic formulae. Then, we consider a more complex setting where we exploit our foundational study of adaptable processes to devise update mechanisms that involve several participants of a distributed protocol, expressed, for example, by a choreography.

In order to introduce adaptation mechanisms and present a spectrum of topologies/evolvability patterns we consider simple extensions to process calculi. Process calculi aim at describing formally the behavior of concurrent systems. A leading motivation in the development of process calculi has been properly capturing the *dynamic character* of concurrent behavior. In fact, much of the success of the π-calculus [21] can be fairly attributed to the way it departs from CCS [20] so as to describe mobile systems in which communication topologies can change dynamically. Subsequent developments can be explained similarly. For instance, the Ambient calculus [18] builds on π-calculus mobility to describe the dynamics of interaction within boundaries and hierarchies, as required in distributed systems. A commonality in these calculi is that the dynamic behavior of a system is realized through a number of *local changes*, usually formalized by reduction steps. Indeed, while in the π-calculus mobility is enforced by the reconfiguration of individual linkages in the communication topology, in the Ambient calculus spatial

José Proença and Massimo Tivoli (Eds.):
14th International Workshop on Foundations of
Coordination Languages and Self-Adaptive Systems (FOCLASA15)
EPTCS 201, 2015, pp. 1–17, doi:10.4204/EPTCS.201.1

mobility is obtained by individual modifications to the containment relations within the ambient hierarchy. This way, the combined effect of a series of changes at a local level (links, containment relations) suffices to explain dynamic behavior at the global (system) level.

There are, however, interesting forms of dynamic behavior that cannot be satisfactorily described as a combination of local changes, in the above sense. These are behavioral patterns which concern change at the *process* level (i.e., the process as a whole), and describe *process evolution* along time. In general, forms of process evolvability are characterized by an enhanced control/awareness over the current behavior and location of running processes. Crucially, this increased control is central to the *adaptation* capabilities by which processes modify their behavior in response to exceptional circumstances in their environment.

This survey reports about the attempt we did to address these shortcomings. In particular, here we present the contents of [2, 3, 1] at a level of detail that allows us to connect them and to discuss the consequences of relating their machinery (this will also lead to new ideas for future work). In particular, we present, in Section 2, a core calculus of adaptable processes expressing a variety of topologies/evolvability patterns and (un)decidability results of verification problems for them (the technical machinery is taken from [2, 3]). We will then, in Section 3, introduce distributed adaptability of protocols as an extension, with update mechanisms similar to those considered for adaptable processes, of the theory of choreography and orchestration/contracts (the technical machinery is taken from [1]). Finally, in Section 4 we discuss some related work and in Section 5 we provide concluding remarks.

2 A Core Calculus of Adaptable Processes: the $\mathscr{E}$ Calculus

We start by presenting *adaptable processes* and results about verification problems for them. The technical machinery presented in this section is taken from [2, 3] (to which the reader is referred for details).

We introduced in [2] the concept of adaptable processes that have a location and are sensible to actions of *dynamic update* at runtime. While locations are useful to designate and structure processes into hierarchies, dynamic update actions implement a sort of built-in adaptation mechanism. We illustrate this novel concept by introducing $\mathscr{E}$, a core process calculus of adaptable processes. The $\mathscr{E}$ calculus arises as a variant of CCS without restriction and relabeling, and extended with primitive notions of *location* and *dynamic update*. In $\mathscr{E}$, $a[P]$ denotes the adaptable process P located at a. Name a acts as a *transparent* locality: P can evolve on its own but also interact freely with its environment. Localities can be nested, and are sensible to interactions with *update prefixes*. An update prefix $\widetilde{a}\{U\}$ decrees the update of the adaptable process at a with the behavior defined by U, a *context* with zero or more holes, denoted by $\bullet$. The *evolution* of $a[P]$ is realized by its interaction with the update prefix $\widetilde{a}\{U\}$, which leads to the process obtained by replacing every hole $\bullet$ in U by P, denoted $U\langle\langle P\rangle\rangle$.

We consider several variants of $\mathscr{E}$, obtained via two orthogonal characterizations. The first one is *structural*, and distinguishes between *static* and *dynamic* topologies of adaptable processes. The former is enforced by constraining the latter so that the topology of adaptable processes $a[P]$ not to vary along the evolution of the system: they cannot be destroyed nor new ones can appear. The second characterization is *behavioral*, and concerns *update patterns*—the context U in an update prefix $\widetilde{a}\{U\}$. As hinted at above, update patterns determine the behavior of running processes after an update action. In order to account for different evolvability patterns, we consider three kinds of update patterns, which determine three families of $\mathscr{E}$ calculi—denoted by the superscripts 1, 2, and 3, respectively. In our view, these variants capture a fairly ample spectrum of scenarios that arise in the joint analysis of correctness and adaptation concerns in evolvable systems. They borrow inspiration from existing programming

languages, development frameworks, and component models.

We now present the $\mathscr{E}$ calculus, its different variants, and its operational semantics. We refer to [2] for further details and discussions. The $\mathscr{E}$ calculus is a variant of CCS [20] without restriction and relabeling, and extended with constructs for evolvability. As in CCS, in $\mathscr{E}$, processes can perform actions or synchronize on them. We presuppose a countable set $\mathscr{N}$ of names, ranged over by $a,b,\ldots$, possibly decorated as $\overline{a},\overline{b}$ and $\widetilde{a},\widetilde{b}$. As customary, we use a and $\overline{a}$ to denote atomic input and output actions, respectively. The syntax of $\mathscr{E}$ processes extends that of CCS with primitive notions of *adaptable processes* $a[P]$ and *update prefixes* $\widetilde{a}\{U\}$.

Definition 2.1 The classes of $\mathscr{E}$ *processes*, *prefixes*, and *update patterns* are described by the following grammars:

$$\begin{array}{rcl} P & ::= & a[P] \mid P \parallel P \mid !\pi.P \mid \sum_{i\in I}\pi_i.P \qquad \pi ::= a \mid \overline{a} \mid \widetilde{a}\{U\} \\ U & ::= & a[U] \mid U \parallel U \mid !\pi.U \mid \sum_{i\in I}\pi_i.U \mid \bullet \end{array}$$

Intuitively, update patterns above represent a context, i.e., a process with zero or more *holes*. The intention is that when an update prefix $\widetilde{a}\{U\}$ is able to interact, the current state of an adaptable process named a is used to fill the holes in the update pattern U. Given a process P, process $a[P]$ denotes the adaptable process P *located at* a. Notice that a acts as a *transparent* locality: process P can evolve on its own, and interact freely with external processes. Localities can be nested, so as to form suitable hierarchies of adaptable processes. The rest of the syntax follows standard lines. A process $\pi.P$ performs prefix π and then behaves as P. Parallel composition $P \parallel Q$ decrees the concurrent execution of P and Q. We abbreviate $P_1 \parallel \cdots \parallel P_n$ as $\prod_{i=1}^{n} P_i$, and use $\prod^k P$ to denote the parallel composition of k instances of process P. Given an index set $I = \{1,..,n\}$, the guarded sum $\sum_{i\in I}\pi_i.P_i$ represents an exclusive choice over $\pi_1.P_1,\ldots,\pi_n.P_n$. As usual, we write $\pi_1.P_1 + \pi_2.P_2$ if $|I| = 2$, and 0 if I is empty. Process $!\,\pi.P$ defines guarded replication, i.e., unboundedly many occurrences of P in parallel, which are triggered by prefix π.

Given an update pattern U and a process Q, we write $U\langle\langle Q\rangle\rangle$ for the process obtained by filling in with Q those holes in U not occurring inside update prefixes (a formal definition can be found in [2]). Hence, $\{\cdot\}$ can be seen as a scope delimiter for holes $\bullet$ in $\widetilde{a}\{U\}$.

We now move on to consider three concrete instances of update patterns U.

Definition 2.2 We shall consider the following three instances of update patterns for $\mathscr{E}$:

1. **Full $\mathscr{E}$ ($\mathscr{E}^1$).** The first update pattern admits all kinds of contexts for update prefixes. This variant, corresponding to the above $\mathscr{E}$ is denoted also with $\mathscr{E}^1$.
2. **Unguarded $\mathscr{E}$ ($\mathscr{E}^2$).** In the second update pattern, holes cannot occur in the scope of prefixes in U:
$$U ::= P \mid a[U] \mid U \parallel U \mid \bullet$$
The variant of $\mathscr{E}$ that adopts this update pattern is denoted $\mathscr{E}^2$.
3. **Preserving $\mathscr{E}$ ($\mathscr{E}^3$).** In the third update pattern, the current state of the adaptable process is always preserved (i.e. "$\bullet$" must occur exactly once in U). Hence, it is only possible to add new adaptable processes and/or behaviors in parallel or to relocate it:
$$U ::= a[U] \mid U \parallel P \mid \bullet$$
The variant of $\mathscr{E}$ that adopts this update pattern is denoted $\mathscr{E}^3$.

$$\text{(SUM)}\quad \sum_{i\in I}\pi_i.P_i \xrightarrow{\pi_j} P_j \ \ (j\in I) \qquad \text{(REPL)}\quad !\pi.P \xrightarrow{\pi} P \parallel !\pi.P \qquad \text{(COMP)}\quad a[P] \xrightarrow{a[P]} \star \qquad \text{(LOC)}\quad \frac{P \xrightarrow{\alpha} P'}{a[P] \xrightarrow{\alpha} a[P']}$$

$$\text{(ACT1)}\quad \frac{P_1 \xrightarrow{\alpha} P_1'}{P_1 \parallel P_2 \xrightarrow{\alpha} P_1' \parallel P_2} \qquad \text{(TAU1)}\quad \frac{P_1 \xrightarrow{a} P_1' \qquad P_2 \xrightarrow{\overline{a}} P_2'}{P_1 \parallel P_2 \xrightarrow{\tau} P_1' \parallel P_2'} \qquad \text{(TAU3)}\quad \frac{P_1 \xrightarrow{a[Q]} P_1' \qquad P_2 \xrightarrow{\widetilde{a}\{U\}} P_2'}{P_1 \parallel P_2 \xrightarrow{\tau} P_1'\{U\langle\langle Q\rangle\rangle/\star\} \parallel P_2'}$$

Figure 1: LTS for $\mathscr{E}$. Rules (ACT2), (TAU2), and (TAU4)—the symmetric counterparts of (ACT1), (TAU1), and (TAU3)—have been omitted.

The process semantics is given in terms of a Labeled Transition System (LTS). It is generated by the set of rules in Figure 1. In addition to the standard CCS actions (input, output, τ), we consider two complementary actions for process update: $\widetilde{a}\{U\}$ and $a[P]$. The former represents the possibility to enact an update pattern U for the adaptable process at a; the latter says that an adaptable process at a, with current state P, can be updated. We define $\longrightarrow$ as $\xrightarrow{\tau}$, and write $P \xrightarrow{\alpha}$ if $P \xrightarrow{\alpha} P'$, for some P'.

Definition 2.3 The LTS for $\mathscr{E}$, denoted $\xrightarrow{\alpha}$, is defined by the rules in Figure 1, with transition labels defined as:

$$\alpha ::= \ a \mid \overline{a} \mid a[P] \mid \widetilde{a}\{U\} \mid \tau$$

In Figure 1, rules (SUM), (REPL), (ACT1), and (TAU1) are standard. Rule (COMP) represents the contribution of a process at a in an update operation; we use $\star$ to denote a unique placeholder. Rule (LOC) formalizes transparency of localities. Rule (TAU3) formalizes process evolvability. To realize the evolution of an adaptable process at a, it requires: (i) a process Q—which represents its current state; (ii) an update action offering an update pattern U for updating the process at a—which is represented in P_1' by $\star$ (cf. rule (COMP)) As a result, $\star$ in P_1' is replaced with process $U\langle\langle Q\rangle\rangle$. Notice that this means that the locality being updated is discarded unless it is re-created by $U\langle\langle Q\rangle\rangle$.

We introduce some definitions that will be useful in the following. We denote with $\rightarrow^*$ the reflexive and transitive closure of the relation $\rightarrow$. We define $Pred(s)$ as the set $\{s' \in S \mid s' \rightarrow s\}$ of *immediate predecessors* of s, while $Pred^*(s)$ denotes the set $\{s \in S \mid s' \rightarrow^* s\}$ of *predecessors* of s. We will also assume point-wise extensions of such definitions to sets, i.e. $Pred(S) = \bigcup_{s\in S} Pred(s)$ and similarly for $Pred^*(S)$.

2.1 The static $\mathscr{E}$ Calculus

Static variants of $\mathscr{E}$ are defined in [2]. Informally speaking, the static characterization of $\mathscr{E}$ processes is related to tree-like structures obtained by nesting of located processes. In *dynamic* adaptable processes, the class here considered, update actions which modify the nesting structure are allowed; in contrast, in static adaptable processes such actions are disallowed: this guarantees that no adaptable process is created nor destroyed along computation. By adding such a static topology constraint we get three static variants $\mathscr{E}_s^1$, $\mathscr{E}_s^2$ and $\mathscr{E}_s^3$ for the three kind of update patterns considered in the $\mathscr{E}^1$, $\mathscr{E}^2$ and $\mathscr{E}^3$ languages. From now on we will denote the latter with $\mathscr{E}_d^1$, $\mathscr{E}_d^2$ and $\mathscr{E}_d^3$ to distinguish them from their static variants.

	$\mathscr{E}_d$ – Dynamic Topology	$\mathscr{E}_s$ – Static Topology
$\mathscr{E}^1$	BA undec / EA undec	BA undec / EA undec
$\mathscr{E}^2$	BA dec / EA undec	BA dec / EA undec
$\mathscr{E}^3$	BA dec / EA undec	BA dec / EA dec

Table 1: Summary of (un)decidability results for dialects of $\mathscr{E}$.

2.2 Eventual and Bounded Adaptation

In [2], we study two *verification problems* associated to $\mathscr{E}$ processes and their (un)decidability. They are defined in terms of standard observability predicates (*barbs*), which indicate the presence of a designated error signal. We thus distinguish between *correct states* (i.e., states in which no error barbs are observable) and *error states* (i.e., states exhibiting error barbs). The first verification problem, *bounded adaptation* (abbreviated BA) ensures that, given a finite k, at most k consecutive error states can arise in computations of the system—including those reachable as a result of dynamic updates. In other words, the number of consecutive erroneous states that can be traversed during a computation is bound by some given number k. The second one, *eventual adaptation* (abbreviated EA), is similar but weaker: it ensures that if the system enters into an error state then it will eventually reach a correct state (the system cannot enter a sequence of consecutive erroneous states that lasts forever). We believe that BA and EA fit well in the kind of correctness analysis that is required in a number of emerging applications. For instance, on the provider side of a cloud computing application, these properties allow to check whether a client is able to assemble faulty systems via the aggregation of the provided services and the possible subsequent updates. On the client side, it is possible to carry out forms of *traceability analysis*, so as to prove that if the system exhibits an incorrect behavior, then it follows from a bug in the provider's infrastructure and not from the initial aggregation and dynamic updates provided by the client.

The main technical results of the paper are summarized in Table 1. The calculus $\mathscr{E}^1$ is shown to be Turing complete, and both BA and EA are shown to be *undecidable* for $\mathscr{E}^1$ processes. The Turing completeness of $\mathscr{E}^1$ says much on the expressive power of update actions. In fact, it is known that fragments of CCS without restriction can be translated into finite Petri nets, so they are not Turing complete. Update actions in $\mathscr{E}$ thus allow to "jump" from finite Petri nets to a Turing complete model. We show that in $\mathscr{E}^2$ BA is *decidable*, while EA remains *undecidable*. Interestingly, EA is already undecidable in $\mathscr{E}^3_d$, while it is *decidable* in $\mathscr{E}^3_s$.

2.3 A Logic for Adaptable Processes

We also considered, in [3], a simple yet expressive temporal logic $\mathscr{L}$ over adaptable processes. Concerning model-checking with formulas of the entire logic $\mathscr{L}$ we have the same decidability results as for the eventual adaptation property, while for a fragment of the logic, denoted with $\mathscr{L}_r$, we have the same decidability results as for the bounded adaptation property (see Table 1).

Definition 2.4 The set At of atomic predicates p is given by the following syntax:

$$\mathrm{p} ::= a \mid \overline{a} \mid T.$$

Predicates a and $\overline{a}$ hold true for states/terms that may perform transitions a and $\overline{a}$, respectively. The intention is that the interpretation of atomic predicates should coincide with the notion of *barb* in the process model. T is the true predicate that holds true for every state/term. In the following, we use α to range over labels $a, \overline{a}$, for some name a.

Definition 2.5 The set $\mathscr{L}$ of logic formulae $\phi, \psi, \ldots$ is given by the following syntax, where $\mathrm{p} \in \mathrm{At}$:

$$\phi ::= \mathrm{p} \mid \phi \vee \phi \mid \phi \wedge \phi \mid \neg\phi \mid \Diamond\phi \mid \Diamond^*\phi$$

The set of logical operators includes atomic predicates $\mathrm{p} \in \mathrm{At}$, the usual boolean connectives ($\vee$, $\wedge$, and $\neg$), as well as dynamic connectives (the next and eventuality modalities, $\Diamond$ and $\Diamond^*$). The interpretation of $\mathscr{L}$ over LTSs is given below, where each formula is mapped into the set of states/terms satisfying it.

$$[\![\alpha]\!] = \{ s \in \mathscr{E} \mid s \xrightarrow{\alpha} \} \qquad [\![T]\!] = \mathscr{E} \qquad [\![\Diamond\phi]\!] = \mathit{Pred}([\![\phi]\!]) \qquad [\![\Diamond^*\phi]\!] = \mathit{Pred}^*([\![\phi]\!])$$
$$[\![\phi_1 \vee \phi_2]\!] = [\![\phi_1]\!] \cup [\![\phi_2]\!] \qquad [\![\phi_1 \wedge \phi_2]\!] = [\![\phi_1]\!] \cap [\![\phi_2]\!] \qquad [\![\neg\phi]\!] = \mathscr{E} \setminus [\![\phi]\!]$$

Connectives are interpreted as usual. We usually write $s \models \phi$ if $s \in [\![\phi]\!]$.

Definition 2.6 A formula ϕ is called *monotone* if it does not contain occurrences of $\neg$.

Restricted monotone formulae are those monotone formulae in which conjunctions are always of the form $\mathrm{p} \wedge \phi$, for some $\mathrm{p} \in \mathrm{At}$ and a monotone formula $\phi \in \mathscr{L}$.

Definition 2.7 A formula ϕ is *restricted monotone* if it is monotone and, for any occurrence of $\phi_1 \wedge \phi_2$ inside ϕ, there exists $i \in \{1,2\}$ such that ϕ_i is a predicate $\mathrm{p} \in \mathrm{At}$.

Restricted logic $\mathscr{L}_r$ is the logic composed by, possibly negated, restricted monotone formulae.

Definition 2.8 The restricted logic is composed by the set $\mathscr{L}_r$ of formulae of the form ϕ or $\neg\phi$, where ϕ is a *restricted monotone* formula.

We now give some examples of formulas in $\mathscr{L}$ and $\mathscr{L}_r$. We take $\Diamond^+\phi \stackrel{\mathsf{def}}{=} \Diamond\Diamond^*\phi$. Possibly the most natural safety property one would like to ensure is the absence of k *consecutive barbs* (representing, e.g., errors, as for the bounded adaptation property):

$$\mathsf{CB}_k(e) \stackrel{\mathsf{def}}{=} \neg\Diamond^*\big(\underbrace{e \wedge \Diamond(e \wedge \Diamond(e \wedge \ldots \wedge \Diamond e))}_{e \text{ appears } k \text{ times}}\big)$$

Observe how $\mathsf{CB}_k(e) \in \mathscr{L}_r$. Another example of a sensible property to ensure is *monotone correctness*: once solved, errors do not reappear. In $\mathscr{L}$ this can be expressed as:

$$\mathsf{MC}(e) \stackrel{\mathsf{def}}{=} \neg\Diamond^*\big(e \wedge \Diamond^+(\neg e \wedge \Diamond^+ e)\big)$$

Assuming a designated barb *ok*, signaling a correct (error-less) state, the above can be captured in $\mathscr{L}_r$ as follows:

$$\mathsf{MC}^r(ok, e) \stackrel{\mathsf{def}}{=} \neg\Diamond^*\big(e \wedge \Diamond^+(ok \wedge \Diamond^* e)\big)$$

The extension of $\mathsf{MC}^r(ok, e)$ to consider k different error phases (it cannot happen that an error re-appears up to k times) is straightforward:

$$\mathsf{MC}^r_k(ok, e) \stackrel{\mathsf{def}}{=} \neg\Diamond^*\big(\underbrace{e \wedge \Diamond^+(ok \wedge \Diamond^*(e \wedge \Diamond^+(ok \wedge \ldots (ok \wedge \Diamond^* e))))}_{ok \text{ appears } k \text{ times}}\big)$$

The latter is used in [3] to model-check a case study about effective scaling in cloud computing, modeled with a $\mathscr{E}^2_d$ process (thus model-checking is decidable).

In [3] we also conjecture that decidability results for $\mathscr{L}_r$ can be extended to monotone formulae (dropping the restricted monotone constraint) and we discuss extensions of the logic with a recursion operator. In particular, we claim that for $\mathscr{L}_r$ with recursion, that would allow us to express properties like eventual adaptation, we have the same (un)decidability results as for formulae of $\mathscr{L}$ (without recursion).

3 Choreographies and Orchestrations with Dynamic Updates

We now discuss a possible approach for expressing distributed dynamic updates of protocol specifications. The technical machinery presented in this section is taken from [7, 1] (to which the reader is referred for details).

Dynamic updates in this context are presented as an extension of our work on choreographies and behavioural contracts in multiparty interactions [5, 7, 9]. We consider distributed adaptability, where a process can update part of a protocol by performing dynamic distributed updates over a set of protocol participants (denoted by *roles* in the choreographical description of the protocol). In particular, we show how update mechanisms considered for adaptable processes can be used to extend the theory of choreography and orchestration/contracts, allowing them to be modified at run-time by internal (self-adaptation) or external intervention. This entails both an extension of choreography and orchestration/contract languages with primitives for distributed updates and named adaptable parts similar to operators $\widetilde{a}\{U\}$ and $a[P]$ considered in the previous section. Moreover also related theories must be extended, such as choreography well-formedness (by, e.g., connectedness constraints) and derivation of participant orchestrations from a choreography by projection.

We start by presenting the choreography and orchestration calculi (representing individual participant behaviour) and then we relate them by projection. Finally, we will extend the the theory presented with distributed dynamic update mechanisms starting from the operators introduced in the previous section.

3.1 The Choreography Calculus

We assume a denumerable set of action names $\mathscr{N}$, ranged over by $a,b,c,\ldots$ and a set of roles *Roles* ranged over by r,s,l.

Definition 3.1 (Choreographies) The set of *Choreographies*, ranged over by $H,L,\cdots$ is defined by the following grammar:

$$H \;::=\; a_{r\to s} \;\mid\; H+H \;\mid\; H;H \;\mid\; H|H \;\mid\; H^*$$

The invocations $a_{r\to s}$ means that role r invokes the operation a provided by the role s. The other operators are choice $\_+\_$, sequential $\_;\_$, parallel $\_|\_$, and repetition $\_^*$.

The operational semantics of choreographies considers two auxiliary terms **1** and **0**. They are used to model the completion of a choreography, which is relevant in the operational modeling of sequential composition. The formal definition is given in Table 2 where we take η to range over the set of labels $\{a_{r\to s} \mid a\in\mathscr{N}, r,s\in \mathit{Roles}\}\cup\{\surd\}$ (the label $\surd$ denotes successful completion). The rules in Table 2 are rather standard for process calculi with sequential composition and without synchronization; in fact, parallel composition simply allows for the interleaving of the actions executed by the operands (apart from completion labels $\surd$ that are required to synchronize).

Choreographies are especially useful to describe the protocols of interactions within a group of collaborating services. To clarify this point, we present a simple example of a protocol described with our choreography calculus.

Example 3.2 (Buyer/Seller/Bank) Let us consider the following choreography composed of three roles: *Buyer, Seller* and *Bank*

$$\begin{array}{l} \mathit{Request}_{\mathit{Buyer}\to\mathit{Seller}};(\ \mathit{Offer}_{\mathit{Seller}\to\mathit{Buyer}}\,|\,\mathit{PayDescr}_{\mathit{Seller}\to\mathit{Bank}}\ );\\ \quad \mathit{Payment}_{\mathit{Buyer}\to\mathit{Bank}};(\ \mathit{Confirm}_{\mathit{Bank}\to\mathit{Seller}}\,|\,\mathit{Receipt}_{\mathit{Bank}\to\mathit{Buyer}}\ ) \end{array}$$

$$a_{r\to s} \xrightarrow{a_{r\to s}} \mathbf{1} \qquad \mathbf{1} \xrightarrow{\surd} \mathbf{0} \qquad H^* \xrightarrow{\surd} \mathbf{0}$$

$$\frac{H \xrightarrow{\eta} H'}{H+L \xrightarrow{\eta} H'} \qquad \frac{H \xrightarrow{\eta} H' \quad \eta \neq \surd}{H;L \xrightarrow{\eta} H';L} \qquad \frac{H \xrightarrow{\surd} H' \quad L \xrightarrow{\eta} L'}{H;L \xrightarrow{\eta} L'}$$

$$\frac{H \xrightarrow{\surd} H' \quad L \xrightarrow{\surd} L'}{H|L \xrightarrow{\surd} H'|L'} \qquad \frac{H \xrightarrow{\eta} H' \quad \eta \neq \surd}{H|L \xrightarrow{\eta} H'|L} \qquad \frac{H \xrightarrow{\eta} H' \quad \eta \neq \surd}{H^* \xrightarrow{\eta} H';H^*}$$

Table 2: Semantic rules for choreographies (symmetric rules omitted).

According to this choreography, the *Buyer* initially sends an item request to the *Seller* that subsequently, in parallel, replies with an offer and sends a payment description to the *Bank*. Then the *Buyer* performs the actual payment to the *Bank*. The latter in parallel replies with a receipt and sends a payment confirmation to the *Seller*.

3.2 The Orchestration Calculus

As for choreographies, we assume a denumerable set of action names $\mathscr{N}$, ranged over by $a, b, c, \ldots$. We use $\tau \notin \mathscr{N}$ to denote an internal (unsynchronizable) computation.

Definition 3.3 (Orchestrations and Systems) The syntax of orchestrations is defined by the following grammar

$$\begin{array}{rcl} C & ::= & \mathbf{0} \mid \mathbf{1} \mid \tau \mid a \mid \overline{a}_r \mid \\ & & C;C \mid C+C \mid C|C \mid C^* \end{array}$$

The set of all orchestrations C is denoted by $\mathscr{P}_{orc}$. Differently from choreographies, orchestration syntax makes it possible to specify orchestrations that fail $\mathbf{0}$ or successfully terminate $\mathbf{1}$. The reason of such a difference is that, conceptually, choreographies represent the desired overall system behaviour, while orchestrations can be used to reason about concrete (possibly failing) participant implementations.
The syntax of systems (parallel composition of orchestrations) is defined by the following grammar

$$P \ ::= \ [C]_r \quad | \quad P \| P$$

A system P is well-formed if: (i) every orchestration subterm $[C]_r$ occurs in P at a different role r and (ii) no output action with destination r is syntactically included inside an orchestration subterm occurring in P at the same role r, i.e. actions $\overline{a}_r$ cannot occur inside a subterm $[C]_r$ of P. The set of all well-formed systems P is denoted by $\mathscr{P}$. In the following we will just consider well-formed systems and, for simplicity, we will call them just systems.

We take α to range over the set of syntactical actions $SAct = \mathscr{N} \cup \{\overline{a}_r \mid a \in \mathscr{N} \wedge r \in Roles\} \cup \{\tau\}$.

The operational semantics of orchestrations is defined by the rules in Table 3 (plus symmetric rules). The operational semantics of systems is defined by the rules in Table 4 plus symmetric rules. We take β to range over the set Act of actions executable by orchestrations and systems, i.e. $Act = SAct \cup \{a_l \mid a \in \mathscr{N} \wedge l \in Roles\} \cup \{a_{r\to s} \mid a \in \mathscr{N} \wedge r,s \in Roles\} \cup \{\overline{a}_{rs} \mid a \in \mathscr{N} \wedge r,s \in Roles\}$. We take λ to range over the set of transition labels $Act \cup \{\surd\}$, where $\surd$ denotes successful termination.

$$\mathbf{1} \xrightarrow{\surd} \mathbf{0} \qquad\qquad \alpha \xrightarrow{\alpha} \mathbf{1}$$

$$\frac{C \xrightarrow{\lambda} C'}{C+D \xrightarrow{\lambda} C'} \qquad \frac{C \xrightarrow{\lambda} C' \quad \lambda \neq \surd}{C;D \xrightarrow{\lambda} C';D} \qquad \frac{C \xrightarrow{\surd} C' \quad D \xrightarrow{\lambda} D'}{C;D \xrightarrow{\lambda} D'}$$

$$\frac{C \xrightarrow{\surd} C' \quad D \xrightarrow{\surd} D'}{C|D \xrightarrow{\surd} C'|D'} \qquad \frac{C \xrightarrow{\lambda} C' \quad \lambda \neq \surd}{C|D \xrightarrow{\lambda} C'|D}$$

$$C^* \xrightarrow{\surd} \mathbf{0} \qquad \frac{C \xrightarrow{\lambda} C' \quad \lambda \neq \surd}{C^* \xrightarrow{\lambda} C';C^*}$$

Table 3: Semantic rules for orchestrations (symmetric rules omitted).

$$\frac{C \xrightarrow{a} C'}{[C]_r \xrightarrow{a_r} [C']_r} \qquad \frac{C \xrightarrow{\overline{a}_s} C'}{[C]_r \xrightarrow{\overline{a}_{rs}} [C']_r} \qquad \frac{P \xrightarrow{\lambda} P' \quad \lambda \neq \surd}{P\|Q \xrightarrow{\lambda} P'\|Q}$$

$$\frac{P \xrightarrow{a_s} P' \quad Q \xrightarrow{\overline{a}_{rs}} Q'}{P\|Q \xrightarrow{a_{r\to s}} P'\|Q'} \qquad \frac{P \xrightarrow{\surd} P' \quad Q \xrightarrow{\surd} Q'}{P\|Q \xrightarrow{\surd} P'\|Q'}$$

Table 4: Semantic rules for systems (symmetric rules omitted).

Here and in the remainder of the paper we use the following notations: $P \xrightarrow{\lambda}$ to mean that there exists P' such that $P \xrightarrow{\lambda} P'$ and, given a sequence of labels $w = \lambda_1\lambda_2\cdots\lambda_{n-1}\lambda_n$ (possibly empty, i.e., $w = \varepsilon$), we use $P \xrightarrow{w} P'$ to denote the sequence of transitions $P \xrightarrow{\lambda_1} P_1 \xrightarrow{\lambda_2} \cdots \xrightarrow{\lambda_{n-1}} P_{n-1} \xrightarrow{\lambda_n} P'$ (in case of $w = \varepsilon$ we have $P' = P$, i.e., $P \xrightarrow{\varepsilon} P$). Moreover, for completely specified systems P (i.e. terms P not included as subterms in a larger term P'), we do not consider transitions corresponding to unmatched input and output actions: namely, we consider only transitions labeled with τ (local internal actions), $\surd$ (global successful termination) and $a_{r\to s}$ (completed interactions).

We now define the notion of correct composition of orchestrations. This notion is the same as in [5]. Intuitively, a system composed of orchestrations is correct if all possible computations may guarantee completion; this means that the system is both deadlock and livelock free (there can be an infinite computation, but given any possible prefix of this infinite computation, it must be possible to extend it to reach a successfully completed computation).

Definition 3.4 (Correct orchestration composition) System $P \in \mathscr{P}$ is a correct orchestration composition, denoted $P\downarrow$, if for every P' such that $P \xrightarrow{w} P'$ there exists P'' such that $P' \xrightarrow{w'} P'' \xrightarrow{\surd}$.

3.3 Choreography Implementation, Projection and Well-Formedness

We are now ready to formalize the notion of correct implementation of a choreography. With $P \xrightarrow{\tau^*} P'$ we denote the existence of a (possibly empty) sequence of τ-labeled transitions starting from the system P and leading to P'. Given the sequence of labels $w = \lambda_1 \cdots \lambda_n$, we write $P \stackrel{w}{\Longrightarrow} P'$ if there exist $P_1, \cdots, P_m$ such that $P \xrightarrow{\tau^*} P_1 \xrightarrow{\lambda_1} P_2 \xrightarrow{\tau^*} \cdots \xrightarrow{\tau^*} P_{m-1} \xrightarrow{\lambda_n} P_m \xrightarrow{\tau^*} P'$.

Intuitively, a system implements a choreography if it is a correct composition of orchestrations and all of its conversations (i.e. the possible sequences of message exchanges), are admitted by the choreography.

Definition 3.5 (Choreography implementation) Given the choreography H and the system P, we say that P implements H (written $P \propto H$) if

- P is a correct orchestration composition and
- given a sequence w of labels of the kind $a_{r\to s}$, if $P \stackrel{w\surd}{\Longrightarrow} P'$ then there exists H' such that $H \xrightarrow{w\surd} H'$.

Note that it is not necessary for an implementation to include all possible conversations admitted by a choreography.

Example 3.6 (Implementation of Buyer/Seller/Bank) As an example, we present a possible implementation of the choreography reported in the Example 3.2.

$$[\overline{Request}_{Seller}; Offer; \overline{Payment}_{Bank}; Receipt]_{Buyer} \parallel$$
$$[Request; (\overline{Offer}_{Buyer} \mid \overline{PayDescr}_{Bank}); Confirm]_{Seller} \parallel$$
$$[PayDescr; Payment; (\overline{Receipt}_{Buyer} \mid \overline{Confirm}_{Seller})]_{Bank}$$

We now present the notion of choreography projection, which yields an orchestration C for each role of a choreography H. The definition is very simple thanks to the idea, we introduced in [7], of projecting communication atoms and then applying homomorphism over all the algebra operators.

Definition 3.7 (Choreography projection) Given a choreography H, the projection H on the role r, denoted with $[\![H]\!]_r$, is defined inductively on the syntax of H in such a way that

$$[\![a_{r\to s}]\!]_t = \begin{cases} \overline{a}_s & \text{if } t = r \\ a & \text{if } t = s \\ \mathbf{1} & \text{otherwise} \end{cases}$$

and that it is a homomorphism with respect to all operators.

It is interesting to observe that given a choreography H, the system obtained composing its projections is not ensured to be an implementation of H. For instance, consider the choreography $a_{r\to s}\ ;\ b_{t\to u}$. The system obtained by projection is $[\overline{a}_s]_r \parallel [a]_s \parallel [\overline{b}_u]_t \parallel [b]_u$. Even if this is a correct composition of orchestrations, it is not an implementation of H because it comprises the conversation $b_{t\to u}a_{r\to s}$ which is not admitted by H.

The problem is not in the definition of the projection, but in the fact that the above choreography cannot be implemented preserving the message exchanges specified by the choreography. In fact, in order to guarantee that the communication between t and u is executed after the communication between r and s, it is necessary to add a further message exchange (for instance between s and r) which is not considered in the choreography.

In order to have the guarantee that the system obtained by projection is consistent with the initial choreography, it is reasonable to consider a subset of *well formed* choreographies. The most general and intuitive notion of well formedness, we introduced in [7], can be obtained by considering only all those choreographies for which the system obtained by projection is ensured to be a correct implementation.

Definition 3.8 (Well formed choreography) A choreography H, defined on the roles $r_1, \cdots, r_n$, is *well formed* if $[\,[\![H]\!]_{r_1}\,]_{r_1} \parallel \cdots \parallel [\,[\![H]\!]_{r_n}\,]_{r_n} \propto H$

It is worthwhile to note that well formedness is decidable. In fact, given a choreography H, it is sufficient to take the corresponding system P obtained by projection, then consider P and H as finite state automata, and finally check whether the language of the first automaton is included in the language of the second one. Note that the terms P and H can be seen as finite state automata thanks to the fact that their infinite behaviours are defined using Kleene-star repetitions instead of general recursion. This decidability result clearly follows from the fact that we restrict to finite state choreographies.

In the literature, syntactic versions of well formedness exist (see e.g. [14, 19, 4]). A sound charcterization of well-formedness is obtained, following [19, 4], by introducing a notion of connectedness for our choreography calculus. The idea is to impose syntactic restrictions in order to avoid the three possible ways in which a system obtained by projection can have a different behaviour w.r.t. its choreography: *connectedness for sequence*, related to the $H;H'$ operator and guaranteeing that an interaction of H' cannot occur before an interaction of H; *unique point of choice*, related to the $H+H'$ operator and guaranteeing that all the roles involved in H and H' are aware of the selected branch; and *no operation interference*, related to the $H|H'$ operator and guaranteeing that a message sent within an interaction of one of H or H' cannot be intercepted by a receive action in an interaction of the other one (see [19, 4] for details).

Theorem 3.9 Let H be a choreography satisfying the *connectedness for sequence*, *unique point of choice* and *no operation interference* conditions. Then H is *well formed*.

3.4 Behavioural Contracts and Contract Refinement

Behavioural contracts make it possible to reason about protocol participant correct composition independently of the language (syntax and semantics) used for expressing the orchestration of a participant. A behavioural contract is defined, basically, as a Labeled Transition System with the same set of labels we considered for the semantics of an orchestration (see, e.g., [10] for a precise definition). An orchestration, therefore, gives rise to a behavioral contract as the labeled transition system obtained by its semantics.

Behavioural contracts are important in the context of service oriented computing for dealing with the problem of service retrieval. Assuming that services expose in their interface an abstract description of their behaviour (a behavioural contract) it is desirable to define an automatic procedure that can be used to check whether a service may correctly play a given role of a choreography. It is therefore crucial to define an, as coarse as possible, notion of *contract refinement* that makes it possible to establish if a discovered service can play a given role, based on the behavioural contract of the service and the behavioural contract derived, by projection, from the choreography. We, thus, define contract refinement as the coarsest preorder over behavioural contracts which preserves correct composition when applied to members of a set of behavioural contracts (one for each role of a choreography). See [5, 6, 7, 8, 9, 10, 11, 12] for formal definitions considering several form of communications/notions of correct composition.

3.5 Introducing Distributed Dynamic Updates

Dynamic updates are specified, similarly as in Section 2, by defining *scopes* $X[H]$ (where X is a scope name) that delimit parts of choreographies that, at runtime, may be replaced by a new choreography, coming from either inside or outside the system. Updates coming from outside may be decided by the user through some adaptation interface, by some manager module, or by the environment. In contrast, updates coming from inside represent self-adaptations, decided by a part of the system towards itself or towards another part of the system, usually as a result of some unsatisfactory interaction. Updates from outside and from inside are indeed quite similar, e.g., an update decided by a manager module may be from inside if the manager behavior is part of the choreography term, from outside if it is not.

Scopes $X[H]$ are updated by means of $\widetilde{X}\{H'\}$ operators similarly as described in Section 2. However, for the sake of simplicity, here we just consider updates $\widetilde{X}\{H'\}$ where H' does not have holes $\bullet$, i.e. in Section 2 this corresponds to the case in which an update pattern U is just a term P (without any $\bullet$). Moreover, we assume $\widetilde{X}\{H'\}$ to just update the content of scope $X[H]$, thus leading to $X[H']$. Therefore, in terms of the machinery in Section 2, this corresponds to having scopes $a[Q]$, for some Q, to be updated by $\widetilde{a}\{U\}$ where U is $a[P]$ for some P, thus $a[Q]$ is replaced by $a[P]\langle\langle Q\rangle\rangle = a[P]$. According to the classification introduced in Section 2, updates of this kind fall in the category of *Unguarded* update patterns (those of $\mathscr{E}_d^2$ language) but not in that of *Preserving* update patterns (those of $\mathscr{E}_d^3$ language). This because in the latter category update patterns are required to have one hole $\bullet$. Moreover, the language topology is not static in that, when a named scope $X[H]$ is updated by means of $\widetilde{X}\{H'\}$, H' can have a different set of named scopes with respect to H and, moreover, due to usage of the ";" operator, named scopes can be created and destroyed at run-time.

Differently from process calculi considered in Section 2 ($\mathscr{E}$ family), here scope updates are not applied by channel based communication (that would not make sense in the context of a choreography, where communication is just a basic action $a_{r\to s}$), but by replacement. Thus $\widetilde{X}\{H'\}$ updates all scopes $X[H]$ (with the same name X) occurring anywhere in the choreography. This also justifies the change of notation from $a[P]/\widetilde{a}\{U\}$ to $X[H]/\widetilde{X}\{H'\}$: in the former a represents a channel on which to communicate, in the latter X represents a syntactic variable to replace. Moreover, notice that, now when a scope $X[H]$ is updated, we are actually performing a *distributed update* that encompasses all the roles involved in H: at the orchestration level, as we will see, we have an X scope for each of such roles and an update on X must update all of them in a distributed way.

Formally speaking, the Choreography Calculus is extended with *scopes* and *updates* as follows (with respect to the notation above we add roles and we omit ˜ in updates):

$$\begin{array}{lll} H ::= & \ldots & \\ \quad | & X : T[H] \qquad (\text{scope}) & \quad | \quad X\{r : H\} \qquad (\text{update}) \end{array}$$

where X is used to range over a set of *scope names* and T is a set of roles.

Construct $X : T[H]$ defines a scope named X currently executing choreography H — the name is needed to designate it as a target for a particular adaptation. Type T is the set of roles (possibly) occurring in the scope. This is needed since a given update can be applied to a scope only if it specifies how all the involved roles are adapted. Operator $X\{r : H\}$ defines *internal updates*, i.e., updates offered by a participant of the choreography. Here r denotes the role offering the update, X is the name of the target scope, and H is the new choreography.

The operational semantics for the new Choreography Calculus with updates is defined only for a proper subset of well defined choreographies.[1] First of all, type T related conditions are added in order

[1] We refer the reader to [1] for a formal definition of this subset of well defined choreographies, here we simply informally

$$(\text{CommUpd})\ \frac{}{X\{r:\ H\} \xrightarrow{X\{r:H\}} \mathbf{1}} \qquad (\text{SeqUpd})\ \frac{H_1 \xrightarrow{X\{r:H\}} H_1'}{H_1;\ H_2 \xrightarrow{X\{r:H\}} H_1';(H_2[H/X])}$$

$$(\text{ParUpd})\ \frac{H_1 \xrightarrow{X\{r:H\}} H_1'}{H_1\ \mid\ H_2 \xrightarrow{X\{r:H\}} H_1'\ \mid\ (H_2[H/X])} \qquad (\text{StarUpd})\ \frac{H_1 \xrightarrow{X\{r:H\}} H_1'}{H_1^* \xrightarrow{X\{r:H\}} H_1';\ (H_1[H/X])^*}$$

$$(\text{ScopeUpd})\ \frac{H_1 \xrightarrow{X\{r:H\}} H_1'}{X:T[H_1] \xrightarrow{X\{r:H\}} X:T[H]} \qquad (\text{ScopeComm})\ \frac{H_1 \xrightarrow{a_{r_1 \to r_2}} H_1'}{X:T[H_1] \xrightarrow{a_{r_1 \to r_2}} X:T[H_1']}$$

$$(\text{Scope})\ \frac{H_1 \xrightarrow{\eta} H_1'}{X:T[H_1] \xrightarrow{\eta} X:T[H_1']}\ \eta \neq X\{r:\ H\} \text{ for any } r,H$$

Table 5: Semantics of Choreographies with updates

to guarantee that: every scope $X:T[H']$ occurring in H has the same type T, denoted with $type(X)$; and every update $X\{r:\ H\}$ injects a choreography H that considers only roles explicitly indicated in the type T of the updated part(s) $X:T[H']$ (at least one scope X is required to occur in the choreography). Moreover, also syntactic restrictions are imposed in order to guarantee that it will never occur that two distinct scopes with the same name are contemporaneously active. This is necessary for the following reason: when a scope is projected, it is distributed among several independent roles running in parallel, and in order to avoid interferences between two distinct scopes with the same name X we assume that only one scope X can be active at a time.

The operational semantics is defined by adding the rules dealing with *scopes* and *updates* reported in Table 5. The transition labels η now include the label $X\{r:\ H\}$ indicating the execution of an update action. In the rules, we use $H[H'/X]$ to denote the substitution that replaces all scopes $X:T[H'']$ with name X occurring in H (not inside update prefixes) with $X:T[H']$.

We briefly comment the rules in Table 5. Rule (CommUpd) makes an internal update available, moving the information to the label. Updates propagate through sequence, parallel composition, and Kleene star using rules (SeqUpd), (ParUpd), and (StarUpd), respectively. Note that, while propagating, the update is applied to the continuation of the sequential composition, to parallel terms, and to the body of Kleene star. Notably, the update is applied to both enabled and non enabled occurrences of the desired scope. Rule (ScopeUpd) allows a scope to update itself (provided that the names coincide), while propagating the update to the rest of the choreography. Rule (Scope) allows a scope to compute.

Example 3.10 (Adaptable Buyer/Seller/Bank) Here, we consider a version of the Buyer/Seller/Bank example discussed in the Example 3.2 where it is possible to update the payment interaction between the buyer and the bank by using, for instance, a new version of the payment protocol according to which the buyer sends its VISA code to the bank and the bank subsequently confirms its correctness. Let us consider the following choreography composed of three roles: *Buyer*, *Seller* and *Bank*

$$\begin{array}{l} Request_{Buyer \to Seller};(\ Offer_{Seller \to Buyer}\,|\,PayDescr_{Seller \to Bank}\ ); \\ \quad X\{Buyer,Bank\}[Payment_{Buyer \to Bank}];(\ Confirm_{Bank \to Seller}\,|\,Receipt_{Bank \to Buyer}\ ) \end{array}$$

report the imposed limitations.

According to the operational semantics defined above, this choreography could, for instance, perform the initial *Request* interaction and then receives an external update:

$$X\{r:\ VISAcode_{Buyer\to Bank}; VISAok_{Bank\to Buyer}\}$$

and then becomes the following choreography:

$$(\ Offer_{Seller\to Buyer} \,|\, PayDescr_{Seller\to Bank}\ );$$
$$X\{Buyer, Bank\}[VISAcode_{Buyer\to Bank}; VISAok_{Bank\to Buyer}]; (\ Confirm_{Bank\to Seller} \,|\, Receipt_{Bank\to Buyer}\ )$$

We are now ready to present the definition of our Orchestration Calculus extended with operators for dynamic updates:

$$\begin{array}{lll} C ::= & ... & \\ \quad | & X[C]^F \qquad (\text{scope}) & \quad | \quad X\{(r_1,\ldots,r_n):\ C_1,\ldots,C_n\} \qquad (\text{update}) \end{array}$$

where F is a *flag* that is either A, standing for active (running) scope, or ε, denoting a scope still to be started (ε is omitted in the following). $X[C]^F$ denotes a scope named X executing C. F is a flag distinguishing scopes whose execution has already begun (A) from scopes which have not started yet (ε). In order for scopes to become active, orchestrations starting execution of scopes with the same name X must synchronize. Also, when all participants in a scope X complete their respective executions, a synchronisation is needed in order to synchronously remove the scope. The update operator $X\{(r_1,\ldots,r_n):\ C_1,\ldots,C_n\}$ provides an update for scope named X, involving roles $r_1,\ldots,r_n$. The new behaviour for role r_i is C_i.

As in the previous sections, systems are of the form $[C]_r$, where r is the name of the role and C its behaviour. Systems, denoted P, are obtained by parallel composition of orchestrations:

$$P \ ::= \quad [C]_r \quad | \quad P \| P$$

In this presentation, we do not formally define a semantics for orchestrations: we just point out that it should include labels corresponding to all the labels of the semantics of choreographies, plus some additional labels corresponding to partial activities, such as an input. We also highlight the fact that all scopes that correspond to the same choreography scope evolve together: they are endowed with *scope start* transitions (transforming a scope from inactive to active, setting its flag F to A) that are synchronized; and with *scope end* transitions (syntactically removing the entire scope) that are synchronized as well. The fact that choreographies feature at most one scope with a given name is instrumental in ensuring this property.

We now discuss how to extend the notion of projection presented in Definition 3.7 for the case without updates.

Definition 3.11 (Projection for choreographies with updates) The projection of a choreography H on a role r, denoted by $[\![H]\!]_r$, is defined as in Definition 3.7 plus the clauses below for scopes and updates:

$$[\![X\{r':\ H\}]\!]_r \ = \ \begin{cases} X\{(r_1,\ldots,r_n):\ [\![H]\!]_{r_1},\ldots,[\![H]\!]_{r_n}\} \text{with } \{r_1,\ldots,r_n\} = type(X) & \text{if } r = r' \\ \mathbf{1} & \text{otherwise} \end{cases}$$

$$[\![X:T[H]]\!]_r \ = \ \begin{cases} X[[\![H]\!]_r] & \text{if } r \in type(X) \\ \mathbf{1} & \text{otherwise} \end{cases}$$

Example 3.12 We now present the projection of the choreography in the Example 3.10 (we omit unnecessary **1** terms):

$$[\overline{Request}_{Seller}; Offer; X[\overline{Payment}_{Bank}]; Receipt]_{Buyer} \parallel$$
$$[Request; (\overline{Offer}_{Buyer} \mid \overline{PayDescr}_{Bank}); Confirm]_{Seller} \parallel$$
$$[PayDescr; X[Payment]; (\overline{Receipt}_{Buyer} \mid \overline{Confirm}_{Seller})]_{Bank}$$

It is interesting to note that the projection clearly identifies where a possible update of the payment should have an effect; namely, only the roles *Buyer* and *Bank* are affected by the update in precise parts of their behaviour. For instance, if $X\{(Buyer, Bank) : (\overline{VISAcode}_{Bank}; VISAok), (VISAcode; \overline{VISAok}_{Buyer})\}$ is executed after the first *Request* interaction occurs, then the system becomes:

$$[Offer; X[\overline{VISAcode}_{Bank}; VISAok]; Receipt]_{Buyer} \parallel$$
$$[(\overline{Offer}_{Buyer} \mid \overline{PayDescr}_{Bank}); Confirm]_{Seller} \parallel$$
$$[PayDescr; X[VISAcode; \overline{VISAok}_{Buyer}]; (\overline{Receipt}_{Buyer} \mid \overline{Confirm}_{Seller})]_{Bank}$$

where the projections of the new protocol are precisely injected in the behaviour of the affected roles.

Ideally, traces of the projected system should correspond to the traces of the original choreography. Actually, we conjecture that this occurs for choreographies satisfying connectedness conditions obtained by extending those already discussed in Section 3.3. We finally point out two main aspects of the expected correspondence result between choreographies and their projections in the case of the calculi extended with dynamic updates. First, labels $X\{r : H\}$ of transitions of the choreography should be mapped to labels of the transitions of the Orchestration Calculus obtained by appropriate label projections. Second, orchestration traces should not consider unmatched input and output labels.

4 Related Work

The $\mathscr{E}$ calculus is related to *higher-order* process calculi such as, e.g., the higher-order π-calculus [22], Kell [23], and Homer [13]. (Further comparisons between $\mathscr{E}$ and other calculi and languages can be found in [2].) In such calculi, processes can be passed around, and so communication involves term instantiation, as in the λ-calculus. Update actions in $\mathscr{E}$ are a form of term instantiation: they can be seen as a streamlined version of the *passivation* operator of Kell and Homer, which allows to suspend a running process. It would be interesting to investigate if the results and techniques developed in this paper can apply to Kell and Homer (or to some interesting fragments of them).

Concerning theories related to choreographies and orchestrations/contracts, among our main contributions we can mention: (i) The formalisation of the relationship between global choreographic descriptions and systems obtained as parallel compositions of peers, (ii) The definition of suitable notions of behavioural contract refinement, and (iii) The proposal of mechanisms for dynamic updates for both the choreography and the orchestration calculi. Concerning (i), we have defined well-formedness for choreographies based on natural projection and the notion of implementation. We have introduced the simple technique of obtaining orchestrations by projection, defining it for communication actions and then exploiting homomorphism over all the process algebraic operators. Moreover our approach leads to a more general notion of well-formedness w.r.t. other approaches like, e.g., [14], where it is defined in terms of three connectedness constraints similar to those we have mentioned in Section 3.3. Concerning (ii), our main contribution is related to the idea of refining all peers guaranteeing that all of them continue to reach local success. This differs from other popular approaches, like those initiated by Fournet et al. [17] or

Padovani et al. [15], where the focus is on the success of one specific peer (usually, the so-called, *client*). Concerning (iii), it is worth to mention that [16] has been a source of inspiration for the present work: the main difference is our choice of expressing adaptation in terms of scopes and code update constructs, rather than using rules. This approach appears more adequate for the definition of a general theory of behavioural typing to be used on more general languages where multiple protocols/choreographies can interleave inside the same program.

5 Conclusion

We just present some observations about the reported results and remarks concerning current/future work.

Choreography and orchestration languages of Section 3 make use, as common in this context, of Kleene-star instead of general recursion (or replication as for $\mathscr{E}$ calculi considered in Section 2). As a consequence (see Section 3.3), they always give rise to finite-state transition systems, for which verification problems are decidable. Given that update mechanisms we introduced in this context belong to the class of those considered in the $\mathscr{E}_d^2$ language of Section 2, verification of properties like bounded adaptability or formulae of a restricted logic like $\mathscr{L}_r$ could be still decidable even if we extend such languages with some more expressive form of recursion. This, for instance, could introduce the possibility, similarly as for the $\mathscr{E}$ family, to generate new participants (or participant instances) at run-time. Notice that such a correspondence between decidability results for the $\mathscr{E}$ family and for choreography and orchestration languages would be possible because, as common in the context of latter languages, name binders (CCS restriction) are not considered, which would otherwise make them Turing complete.

We are currently working on applying the theory of updatable choreographies/orchestrations in the context of session types for typing a concrete language with session spawning, where choreographies play the role of global types attached to sessions and we use orchestrations for checking, via typing rules, that the code actually conforms with the specified global types. In this context, extending our contract refinement theory [5, 7, 9, 10, 12] to updatable choreographies/orchestrations would make it possible to define a notion of semantic subtyping.

References

[1] Bravetti M., Carbone M., Hildebrandt T., Lanese I., Mauro J., Perez J.A., Zavattaro G.: Towards Global and Local Types for Adaptation. *Proc. of* 2nd International Workshop on Behavioural Types – SEFM'13 Colocated Workshops pages 1–12, volume 8378 of *Lecture Notes in Computer Science*, Springer-Verlag, 2013. doi:10.1007/978-3-319-05032-4_1

[2] Bravetti, M., Di Giusto, C., Pérez, J. A., Zavattaro, G.: *"Adaptable processes"*, in *Logical Methods in Computer Science*, 8(4), 2012. doi:10.2168/LMCS-8(4:13)2012

[3] Bravetti, M. , Di Giusto, C., Pérez, J. A., Zavattaro, G.: *"Towards the Verification of Adaptable Processes."*, in Proc. of 5th International Symposium on *Leveraging Applications of Formal Methods, Verification and Validation. Technologies for Mastering Change (ISoLA 2012) Part I*, LNCS 7609: 269–283, Springer 2012. doi:10.1007/978-3-642-34026-0_20

[4] Bravetti, M., Lanese, I., Zavattaro, G.: Contract-Driven Implementation of Choreographies. *Proc. TGC 2008*: LNCS 5474, Springer, 2009, 1–18. doi:10.1007/978-3-642-00945-7_1

[5] Bravetti, M., Zavattaro, G.: Contract based Multi-party Service Composition, *Proc. FSEN'07*, LNCS 4767, Springer, 2007, 207–222. doi:10.1007/978-3-540-75698-9_14

[6] Bravetti, M., Zavattaro, G.: A Theory for Strong Service Compliance, *Proc. Coordination'07*, LNCS 4467, Springer, 2007, 96–112. doi:10.1007/978-3-540-72794-1_6

[7] Bravetti, M., Zavattaro, G.: Towards a Unifying Theory for Choreography Conformance and Contract Compliance, *Proc. SC'07*, LNCS 4829, Springer, 2007, 34–50. doi:10.1007/978-3-540-77351-1_4

[8] Bravetti, M., Zavattaro, G.: Contract Compliance and Choreography Conformance in the Presence of Message Queues, *Proc. WS-FM'08*, volume to appear of *LNCS*, 2008. doi:10.1007/978-3-642-01364-5_3

[9] Bravetti, M., Zavattaro, G.: A Foundational Theory of Contracts for Multi-party Service Composition, *Fundamenta Informaticae* 89(4):451–478, IOS Press, 2008.

[10] Bravetti, M., Zavattaro, G.: Contract-Based Discovery and Composition of Web Services, *9th International School on Formal Methods for the Design of Computer, Communication and Software Systems: Web Services (SFM-09:WS)*, LNCS 5569, Springer, 2009, 34–50. doi:10.1007/978-3-642-01918-0_7

[11] Bravetti, M., Zavattaro, G.: A Theory of Contracts for Strong Service Compliance, *Mathematical Structure in Computer Science* 19(3):601–638, Cambridge University Press, 2009. doi:10.1017/S0960129509007658

[12] Bravetti, M., Zavattaro, G.: Service Discovery and Composition based on Contracts and Choreographic Descriptions, *Adaptive Web Services for Modular and Reusable Software Development: Tactics and Solutions*, IGI Global, 2012. doi:10.4018/978-1-4666-2089-6

[13] M. Bundgaard, J. C. Godskesen, and T. Hildebrandt. Bisimulation congruences for homer — a calculus of higher order mobile embedded resources. Technical Report TR-2004-52, IT University of Copenhagen, 2004.

[14] Carbone M., Honda K., Yoshida N.: Structured Communication-Centred Programming for Web Services. In *ESOP'07*, volume to appear of LNCS, 2007. doi:10.1007/978-3-540-71316-6_2

[15] Carpineti S., Castagna G., Laneve C., Padovani L.: A Formal Account of Contracts for Web Services. In *WS-FM'06*, volume 4184 of LNCS, pages 148-162, 2006. doi:10.1007/11841197_10

[16] Dalla Preda M., Giallorenzo S., Lanese I., Mauro J., Gabbrielli M.: AIOCJ: A Choreographic Framework for Safe Adaptive Distributed Applications. *Proc. SLE 2014*: LNCS 8706, Springer, 2014, 161–170. doi:10.1007/978-3-319-11245-9_9

[17] Fournet C., Hoare C.A.R, Rajamani S.K., Rehof J.: Stuck-Free Conformance. In *Proc. CAV'04*, volume 3114 of LNCS, pages 242–254, 2004. doi:10.1007/978-3-540-27813-9_19

[18] L. Cardelli and A. D. Gordon. Mobile ambients. *Theor. Comput. Sci.*, 240(1):177–213, 2000. doi:10.1016/S0304-3975(99)00231-5

[19] Lanese I., Guidi C., Montesi F., Zavattaro G.: Bridging the gap between Interaction- and Process-Oriented Choreographies. *Proc. of* 6th IEEE International Conferences on Software Engineering and Formal Methods (SEFM'08), pages 323–332, IEEE Computer Society press, 2008. doi:10.1109/SEFM.2008.11

[20] R. Milner. *Comunication and Concurrency*. Prentice Hall, 1989.

[21] R. Milner, J. Parrow, and D. Walker. A Calculus of Mobile Processes, I. *Inf. Comput.*, 100(1):1–40, 1992. doi:10.1016/0890-5401(92)90008-4

[22] D. Sangiorgi. *Expressing Mobility in Process Algebras: First-Order and Higher-Order Paradigms*. PhD thesis CST–99–93, University of Edinburgh, Dept. of Comp. Sci., 1992.

[23] A. Schmitt and J.-B. Stefani. The kell calculus: A family of higher-order distributed process calculi. In *Global Computing*, volume 3267 of *LNCS*, pages 146–178. Springer, 2004. doi:10.1007/978-3-540-31794-4_9