Towards Timed Requirement Verification for Service Choreographies

Nawal Guermouche[1,2]
[1] CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France
[2] Univ de Toulouse, INSA, LAAS, F-31400 Toulouse, France
nguermou@laas.fr

Silvano Dal Zilio[1,2]
[1] CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France
[2] Univ de Toulouse, LAAS, F-31400 Toulouse, France
dalzilio@laas.fr

***Abstract*—In this paper, we propose an approach for analyzing and validating a composition of services with respect to real time properties. We consider services defined using an extension of the Business Process Execution Language (BPEL) where timing constraints can be associated to the execution of an activity or define delays between events. The goal is to check whether a choreography of timed services satisfies given complex real time requirements. Our approach is based on a formal interpretation of timed choreographies in the Fiacre verification language that defines a precise model for the behavior of services and their timed interactions. We also rely on a logic-based language for property definition to formalize complex real-time requirements and on specific tooling for model-checking Fiacre specifications.**

***Keywords*-Timed BPEL processes, choreography analysis, asynchronous services, real-time requirements, formal verification.**

I. Introduction

Web Services are a set of standards that enable the definition of complex software systems based on the composition of autonomous and heterogeneous services. Web Services programming relies on a set of XML based standards, such as the Web Service Description Language (WSDL)[1], for the description of services interface, or the Business Process Execution Language (BPEL)[2], for defining service orchestration. The notion of *choreography* (see e.g. [4]) is useful to reason about the collaboration of services from a global viewpoint. Basically, a choreography provides a way to specify the overall behavior expected from the composition of services. Since time plays a crucial role when reasoning about business processes, we need to be able to express the time related attributes of a choreography and we need to be able to check timing requirements on them.

Taking into account temporal (quantitative) aspects in the specification of services improves expressiveness. However, it makes reasoning about service composition much harder and makes known problems more challenging such as *timed conformance* [15] and *timed compatibility* problems [19]. While timed conformance and compatibility are important, alone they are not sufficient. In fact, in addition to these features, it is crucial to check time related properties of compositions. For instance, it is useful to check the "worst-case execution time" of a composition [24] or to check that a partial deadline is met. Particularly, in the context of our work, we are interested in checking complex time related properties of a choreography, such as verifying that once an event evt_1 arises, if an event evt_2 does not occur within a delay, then an event evt_3 must not occur before a given delay. To do so, we define a formal model for the semantics of timed choreographies and propose a model-based approach for checking complex real-time requirements, which extends our previous work on compatibility analysis [17]. Complex requirements we consider are defined using a logical-based formalism that is able to express real-time constraints between the occurrence of events; we use a property pattern language that defines a fragment of Metric Interval Temporal Logic (MITL), a real-time extension of Linear Temporal Logic. As a result, our framework can be used to check very general (complex) properties, that go beyond the mere absence of deadlocks or the worst-case execution time. We give some examples of real-time requirements that can be checked automatically in Section III.

More precisely, in this paper, we consider *timed services* extended with timing constraints. Our goal is to check whether a *timed choreography*, obtained from the composition of timed services, satisfies given complex requirements. We choose a rich model for expressing timing information where constraints can be associated to the execution time of (basic and structured) activities as well as on the delay between two events. In our context, an event may be *local* to a given service—for instance an invocation or the end of an activity—or may be *global*, associated to a message crossing service boundaries.

The formal semantics of timed choreographies is defined using an interpretation of services in Fiacre [8], [9], a formal modeling language that can be used to represent the behavioral and timing aspects of a system. The use of Fiacre is well adapted for this context since it captures efficiently timing and concurrency aspects of systems. It has been designed both as the target language of model transformation engines—interpretation have been defined for system description languages such as SDL, UML or

[1] http://www.w3.org/TR/WSDL/
[2] http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0.html

AADL—and as the source language for formal verification toolboxes, such as CADP or Tina [10], the TIme Petri Net Analyzer. In Section V we show how we can use our interpretation of timed BPEL processes in Fiacre and the Tina model-checking tools in order to check complex real-time requirements of timed choreographies.

The paper is organized as follows. In Section II we describe related works in the domain of the formal analysis of Web Services and outline our contributions. In Section III we define an example of timed choreography that is used to give an informal overview of our framework. Section IV presents our formal framework. Before concluding, we report on some experimental results obtained using a prototype implementation of our framework.

II. Related Work

We list a series of works related to the formal analysis of Web Services. In [13], a Web service (peer) is modeled as non-deterministic communicating finite state machine. In this work, a choreography of services is specified as a collaboration diagram. In order to verify the reliability of a choreography, the authors propose to check if the peers realize the collaboration diagram. In this work, the authors are not interested in timed properties, which are very important. In [22], Hwang et al focus on the problem of verifying the conformance of a set of Web services to a given choreography model. This problem consists in checking whether a set of Web services respects the specified choreography interaction model. As the work presented in [13], Hwang et al do not cater for timed properties.

In [6], [7], [11], the authors address the problem of checking compatibility between two services. Their approach is restricted to synchronous services and is exclusively based on the sequences of messages that can be exchanged by a service. These assumptions are quite restrictive since, for instance, two services may engage in a successful conversation even if their behaviors do not have the same branching structure. On the other hand, compatible services may exhibit very different behaviors when time is taken into account. Compatibility between timed services has been first studied by Benatallah et al. [5]. This work has been latter extended in [28]. In these approaches, it is only possible to specify the delay between two messages inside the same service, whereas we consider richer time constraints. Moreover, these works address only one specific requirement—absence of deadlocks—while we allow the verification of more general complex properties.

Eder and Tahamtan introduce the notion of *time conformance* in [15]. This is defined as the problem of checking whether a timed orchestration satisfies a given timed choreography. In this framework, time constraints are limited to expressing the execution time of basic activities. This approach is different in nature from the one followed in our work. Indeed, our goal is to check properties of choreographies; we do not suppose that a given choreography "is correct" and that parts of its implementation—a timed orchestration—should conform to this specification. We should point out that, in our case, we solve a model-checking problem for a real-time extension of temporal logic, since we consider timed models, and for a dense time model. Another instance of a "behavior conformance" problem is studied by Ham et al [21], [20] that focuses on the problem of substituting a service by another. These works address both the temporal and financial costs of operations. The authors do not consider time constraints over structured activities or across service boundaries and are not interested in checking complex properties of timed choreographies.

Kazhamiakin et al. [24] adopt a formalism closely related to timed automata to model the behavior of a timed orchestration. This work is based on a discrete-time variant of the Duration Calculus for expressing requirements, while we work with a dense model of time. Furthermore, they consider—other than timeouts—the temporal cost of manual activities and deal with synchronous services, while, in our framework, we propose a richer temporal model. Another difference is that we allow the declaration of more complex and expressive real-time requirements, like those related to the *absence pattern*, that states that parts of a system's execution are free of some events (e.g., once an event evt_1 arises, if an event evt_2 does not occur within a delay, so an evt_3 must arise after a duration).

In [23], Kallel et al adopt the formalism of XTUS-Automata to formalize the behaviour of services extended with timed constraints related to relative and absolute time. Based on this model, the authors present a verification process which aims at checking absence of deadlocks. In the same context and using Petri net theory, in [26], the authors address the problem of assigning and verifying conformance constraints of deadlines. Although these approaches consider timed properties, they do not allow checking complex timed requirements.

Our approach extends previous works based on timed automata [18], [17] that are concerned with the compatibility analysis of timed services. In this work, verification is restricted to checking the absence of deadlock in a composition of services. Moreover, the model can only express temporal constraints on the delays between the exchange of messages and not on the duration of activities.

Our Contributions. Our first contribution lies in the definition of a rich model for timed services. This model integrates efficiently the real-time and concurrent characteristics of timed services and is compatible with both synchronous and asynchronous services.

Our second contribution consists in defining new class of real-time requirements that goes far beyond the mere absence of deadlock or the worst-case execution time. In Section III we illustrate the use of this model in the context of a healthcare scenario and show how time constraints and

requirements can be expressed. We also explain where this information could be reasonably stored: WSDL and BPEL for local constraints and Service Level Agreement (SLA) contracts for global constraints.

Our final contribution is the definition of a set of transformation rules of timed business processes into Fiacre, a formal verification language with native support for expressing concurrency and time-constrained interactions. By using the result of the transformation, we apply model-checking primitives in order to automatically check complex real time requirements. The proposed approach has been implemented in a tool that takes a collection of annotated BPEL processes and a set of real-time requirements as input, and as a result it validates the set of specified requirements.

III. Global Overview of the Modeling Framework

This work is part of the National Project ITEMIS[3] and the European JU Artemisia CESAR[4] project.

ITEMIS aims at defining a reference architecture, methodologies, and a set of techniques and tools for the verification and systems development. CESAR aims at providing techniques and methodologies to develop reliable embedded systems. To do so, multi-viewpoint based development processes must be considered to tackle not only functional aspects but also safety, costs, robustness, timed constraints, etc. These properties must be captured and formalized to allow validation and verification to be performed [1]. In our work, we tackle the problem of verifying complex timed properties of service based systems.

Next, we will present the underlying features of our framework with the help of an example of timed choreography. Our modeling framework relies on a syntax for expressing services (based on BPMN); temporal constraints annotations; and a requirement language based on real-time property specification patterns. The verification framework is defined in the next section. We opted for a scenario from the healthcare domain related to patient handling during a medical examination. The scenario involves three entities, each one managed by a service: (1) the medical consultation clinic; (2) a medical analysis laboratory; and (3) a pharmacy.

The choreography is depicted in Fig. 1, where each service is in a different swimlane. With the aid of the *medicalConsultationService* (MCS), a doctor can check the social security number (ssn) of a patient. If the ssn is valid, then the MCS may ask the *medicalAnalysisService* (MAS) to perform some medical analyses and, in parallel, asks for a radiography. Once the medical analyses are fulfilled—and after analyzing the different medical data—a medical report is compiled and drugs can be ordered.

Temporal Constraints

The services may interact using asynchronous and synchronous exchange of messages. As usual, the choreography imposes constraints on the order of these messages. It also defines temporal constraints through the use of annotations. We consider two kinds of temporal constraints in our framework, local and global.

Local temporal constraints: are associated to the execution of a service. They are used to specify the duration and/or the delays required to perform an activity. We consider two kinds of local constraints:

(1) *temporal costs* are used to define the estimated execution time of an activity, say A, and are of the form $d(A) \in I$, where I is a time interval[5]. For instance, the medical report operation in the MCS requires at least 2 hours: $d(medicalReport) \in [2, ..[$. The most suitable place to store these constraints is in the WSDL file that describes the operations;

(2) *temporal delays* are used to specify the expected delay between two activities and are of the form $d(A_1, A_2) \in I$. For instance, the Pharmacy Service performs the *drugsChecking* activity between 6–12 hours after the start of the service: $d(drugsRequest, drugsChecking) \in [6, 12]$. A temporal delay expresses a commitment from the service, much like a timeout; it is not a requirement that should be checked a posteriori. The most suitable place to store temporal delays is to add annotations in the BPEL file that describes the service.

Global temporal constraints: are used to specify the temporal contract of a service and are associated to pair of messages (dotted lines in our diagram) that cross service boundaries. For example, the Pharmacy Service promises that the drug order is delivered within 24–48 hours of the drug request: $d(msg_3, msg_4) \in [24; 48]$. Global constraints live at the same level than Service Level Agreement contracts.

Given these services, our goal is to verify if their choreography satisfies a set of real-time requirements, such as:

- Once drugs are requested, if the request is not canceled within 6 hours, then they should not be changed for another 48 hours.
- Two medical analyses are not allowed for a patient in less than 10 days.
- Once the medical examination starts, drugs must be delivered within 48 hours.

In the following section, we will present our framework.

[3]ITEMIS: Integrated information and embedded systems. It is an ANR (National Research Agency) project (2009-2012).

[4]CESAR: Cost-efficient methods and processes for safety relevant embedded systems. It is a European funded project from ARTEMIS JOINT UNDERTAKING (JU) (2009-2012).

[5]Time intervals may be open, like $]2.5, 3]$, or unbounded, like $[1, ..[$.

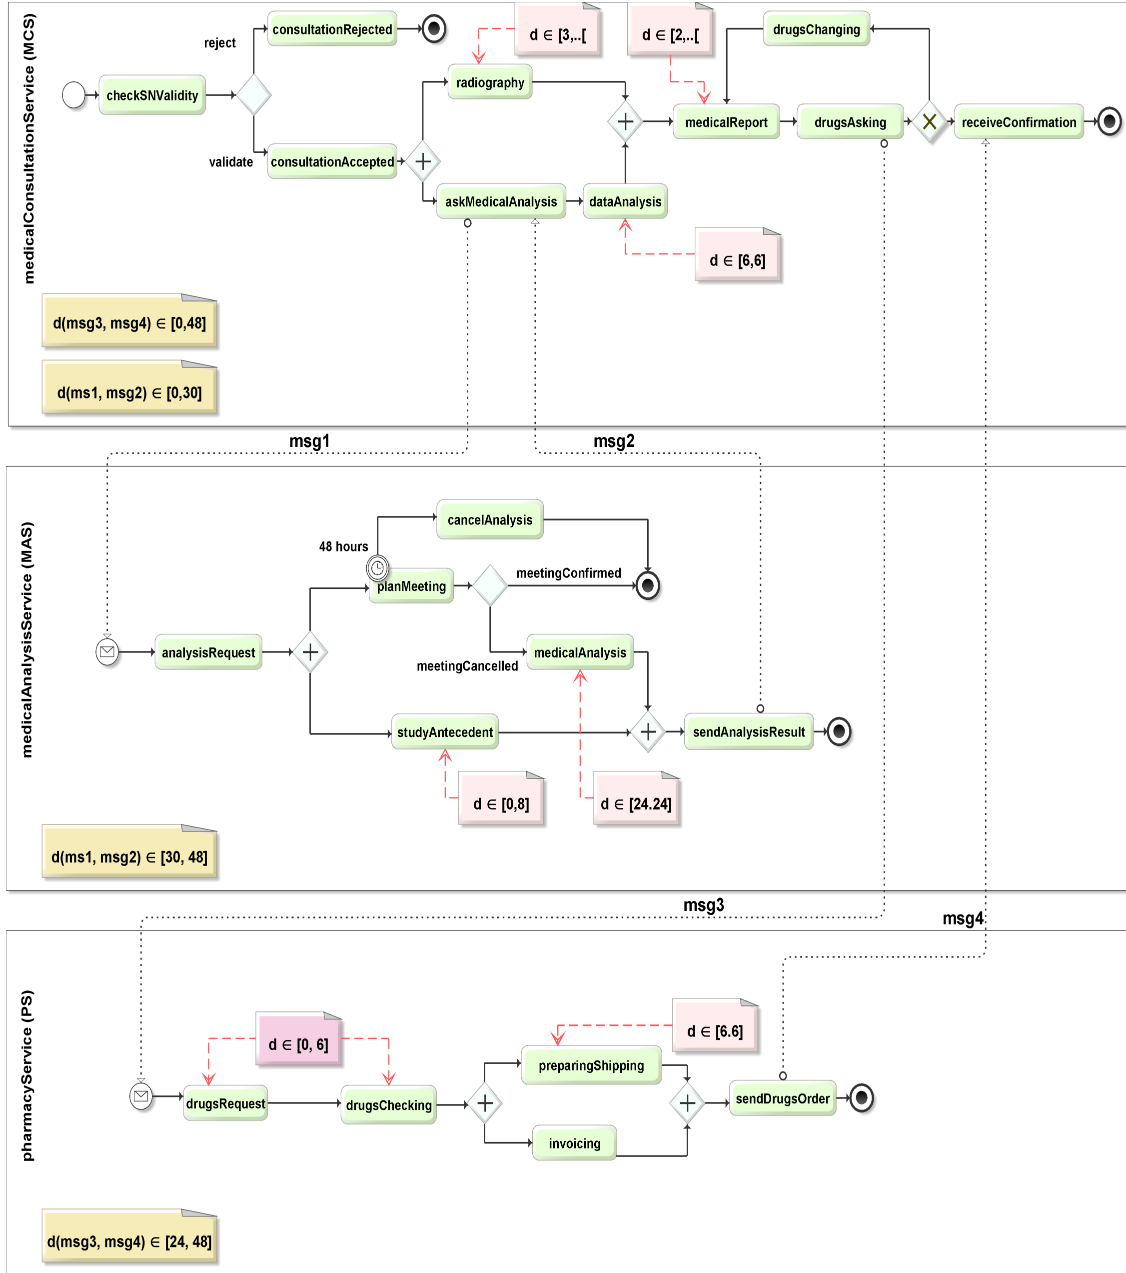

Figure 1. Global view of the health-care application

IV. Verification Framework

In this section, we consider the framework used to define the formal semantics of timed choreographies. Our approach is based on an interpretation of services in Fiacre [8], [9], a formal verification language that can model the behavioral and timing aspects of a system. This method has some advantages compared to related work where the semantics of services is given using a dedicated formalism (see e.g. [25]) or a low-level formalism, like timed automata [18]. Indeed, Fiacre provides high-level operators, special support for different concurrency paradigm and a hierarchical (component-based) structure that simplify the encoding of system description languages. Moreover, the language comes equipped with a set of dedicated tooling: compilers to different model-checking tool suite (like CADP or Tina [10]); support for the real-time requirement language described in the section IV-C; and support for Model-Driven Engineering. In particular, Fiacre is the intermediate language used for model verification in Topcased [8], [16]—an Eclipse based toolkit for critical systems (http://www.topcased.org/)—where it is

used as the target of model transformation engines from various languages, such as SDL, UML or AADL.

A. The Fiacre Language

Fiacre offers a formal representation of both behavioral and timing aspects of systems for formal verification and simulation purposes. The design of the language is inspired from decades of research on concurrency theory and real-time systems theory. For instance, its timing primitives are borrowed from Time Petri nets [27], while the integration of time constraints and priorities into the language can be traced to the BIP framework [12]. Fiacre processes can interact both through synchronization (message-passing) and access to shared variables (shared-memory). A formal definition of the language is given in [9].

Fiacre programs are stratified in two main notions: *processes*, which are well-suited for modeling structured activities, and *components*, which describes a system as a composition of processes, possibly in a hierarchical manner. The language is strongly typed, meaning that type annotations are exploited in order to guarantee the absence of unchecked run-time errors.

A *process* is defined by a set of *control states* (say s_0, s_1, ...) and parameters, each associated with a set of *complex transitions*, which are programs specifying how parameters are updated and which transitions may fire. For example, the process declaration:

```
process P[q : none](&v : nat) is ...
```

expresses that P is a process that uses the port q for synchronization—the port carries no data—and has one parameter, v, that is a (reference to a) shared variable holding natural values. Complex transitions are built from expressions and deterministic constructs available in typical programming languages (assignments, conditionals, while loops and sequential compositions), nondeterministic constructs (nondeterministic choice and assignments) and communication events on ports. For example, the transition:

```
from s0 select v :=v+1; to s1
        []     on(v=0); to s0
        end
```

states that, in state s_0, the process may choose non-deterministically between two alternatives. Either increments the value of the variable v and moves to s_1, or loops to s_0 if v is nil.

A *component* is defined as the parallel composition of processes and/or other components, expressed with the operator **par** ... || ... **end**. While components are the unit of composition, they are also the unit for process instantiation and for ports and shared variables creation. The syntax of components allows to restrict the access mode and visibility of shared variables and ports, to associate timing constraints with communications and to define priority between communication events. For example, the following declaration states that C is a component with a private port r—synchronization over r is in time 0—and two fresh instances of the process P.

```
component C(x : nat) is
    port r:none in [0,0]
    var v1:nat :=x, v2:nat :=3
    par P[r](v1) || P[r](v2) end
```

B. Interpretation of BPEL in Fiacre

In order to perform verification process, we define an interpretation of timed BPEL processes where a service is modeled by a Fiacre component and such that service invocation (both synchronous and asynchronous) is modeled using shared variables. In this approach, the interpretation of a choreography is obtained from the parallel composition of its services. We only detail the encoding of a representative subset of activities. Moreover, we note that our approach does not take into account the expected behaviour in case of faults or communication failures, which is out of scope for this paper.

Interpretation of communication.: We model communication using a shared variable that acts as a buffer counting the number of messages exchanged between services; each WSDL message, say msg_i, is represented by an integer variable, $\mathtt{msgVar}_i$, with initial value 0. Then (an asynchronous) message emission is encoded by incrementing the variable and reception by decrementing it. For example, reception of the message msg is encoded by the Fiacre expression: **on**(msgVar>0); msgVar:=msgVar−1 (we use the **on** expression to test that the "message channel" is not empty). This approach can be extended to take into account the values exchanged in a message, as long as the number of values stay finite.

Interpretation of parallel activities as processes.: A service S is encoded by a component S. We encode each parallel activity A_i in S by a Fiacre process $\mathtt{A}_i$ with two specific ports: ps for signaling the start of the activity and pe for signaling its end. A <flow> activity in BPEL is used to execute sub-activities, $A_1, \ldots, A_n$, in parallel. In our encoding, a flow activity inside the service S is turned into a parallel composition $\mathtt{A}_1$(ps,pe) || ... || $\mathtt{A}_n$(ps,pe), such that instances of the $\mathtt{A}_i$'s are synchronized on their start and end event. Hereafter, we define the processes corresponding respectively to the basic activities <receive>, <reply> and <invoke>

```
process Rcv[ps:none, pe:none] (&msgVar:nat) is
  states start, s1, s2, s3 init to start
  from start ps; to s1
  from s1 on(msgVar>0);wait[0,0];
       msgVar:=msgVar-1;to s2
  from s2 pe; to s3
```

```
process Rpl[ps:none, pe:none](&msgVar:nat) is
  states start, s1, s2, s3 init to start
  from start ps; to s1
  from s1 wait[0,0];msgVar:=msgVar+1;to s2
  from s2 pe; to s3
```

```
process Invk[ps:none, pe:none](&msgOutVar,
        &msgInVar:nat)is
  states start, s1, s2, s3, s4 init to start
  from start ps; to s1
  from s1 wait[0,0];msgOutVar:=msgOutVar+1;to s2
  from s2 on(msgInVar>0);wait[0,0];
          msgInVar:=msgInVar-1; to s3
  from s3 pe; to s4
```

Interpretation of sequential basic activities.: Basic activities that are not executed in parallel are mapped to a state in a single Fiacre process, say P. The goal is to obtain a more efficient encoding (with less processes and states). In this context, an activity A_i is mapped to a state s_i in P and its effect is implemented by a transition from the state s_i to the state s_j such that A_j is the next in line from A_i in the sequence.

<receive>

```
from si on (MsgVar>0); wait[0,0];
        msgVar:=msgVar-1; to sj
```

<reply>

```
from si wait[0,0]; msgVar:=msgVar+1; to sj
```

```
<invoke>

from si wait[0,0]; msgOutVar:=msgOutVar+1;
        to si'
from si' on(msgInVar>0); wait[0,0];
        msgInVar:=msgInVar-1; to sj
```

Interpretation of temporal constraints.: We start by describing the interpretation of temporal cost. An activity that can be executed within a delay is modeled as a transition that takes place after a timed interval. We consider two categories of activity that can have a temporal cost: *one-way operation* and *request-response operation*. In this context, we map an operation A_i to a pair of states s_i and s_i'. The effect of A_i is implemented by a sequence of two transitions from the state s_i to the state s_j, where A_j is the activity that logically follows A_i. We give below the encoding of one-way and request-response <invoke> operations with an execution time of d.

(one-way)

```
from si on(msgVar>0);wait[0,0];msgVar:=msgVar-1;
        to si'
from si' wait[d,d];to sj
```

(request-response)

```
from si on(msgVar>0);wait[0,0];msgVar:=msgVar-1;
        to si'
from si' wait[d,d];msgVar:=msgVar+1; to sj
```

The interpretation of temporal delays and global constraints are similar. We associate to every delay constraint of the form $d(A_1, A_2) \in I$ a process in Fiacre, say TObs. The role of this process is to observe the delay between the end of A_1 (synchronization on the port pe_1) and the start of A_2 (synchronization on ps_2). The encoding is very similar for a delay between messages. We give below the "time observer" process corresponding to a constraint of the form $d(A_1, A_2) \in [0; d]$ (the **unless** operator is used to state that the transition to err has an higher priority). We can test if the constraint is violated by checking whether the process TObs can reach the state err.

```
process TObs[pe1, ps2:none]() is
  states start, s1, s2, err init to start
  from start pe1; to s1
  from s1 select ps2;to s2
          unless wait[d,d];to err
          end
```

Interpretation of timeouts (onAlarm).: We describe our interpretation of a timer-based alarm, <onAlarm for ="d">, using the example of a simple <pick> activity A_i such that:

```
Ai = <pick><onMessage name="m"/>
       <onAlarm for="d">Ak</onAlarm>
     </pick>
```

meaning that the activity will select the activity A_k after d unit of time unless it receives the message m before. The activity A_i may be encoded using the following transition:

```
from si select wait[0,0];on(msgVar>0);
               msgVar:=msgVar-1;to sj
        unless wait[d,d];to sk
        end
```

Example 1: (Interpretation of Pharmacy Service)

We take the example of the healthcare scenario introduced in section III to show our interpretation. The architecture of the Fiacre component corresponding to the Pharmacy Service is displayed in Fig. 2.

The Pharmacy Service (PS) has three temporal constraints: (c_1) the activity *drugsChecking* must be done within 6–12 hours of receiving the drugs request (a local, delay constraint); (c_2) the execution time of activity *preparingShipping* is 6 hours (a local, temporal constraint); and (c_3) the message sent by activity *sendDrugsOrder* should be emitted within 24–48 hours from receiving the drugs request (this is a global constraint).

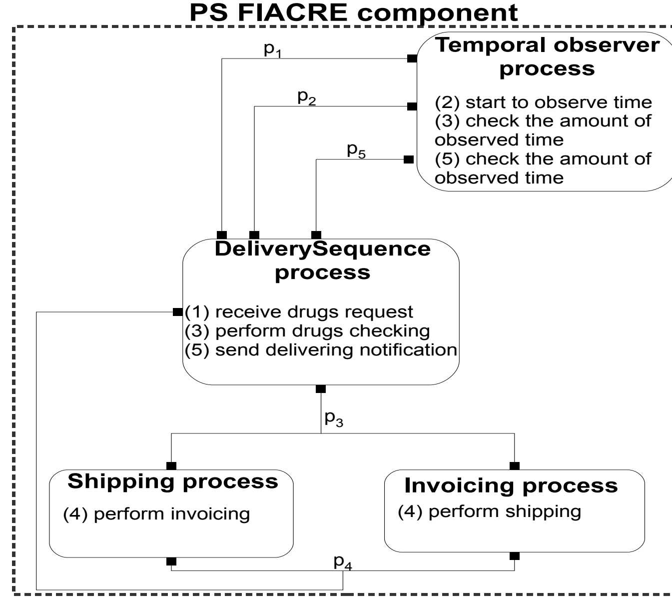

Figure 2. Architecture of the Fiacre process for the Pharmacy Service

The PS is built from a <sequence> activity that contains two operations that should be performed in parallel (inside a <flow> activity): *preparingShipping* and *invoicing*. Therefore, in our (optimized) interpretation, the component PS contains the parallel composition of three processes— one for the sequence and one for each activity inside the <flow>— added to a "time observer" process for the two constraints c_1 and c_3. These four processes appear in the diagram of Fig. 2 with an explicit naming convention.

We can describe the possible interactions of the PS component as follows. First, the deliverySequence process starts by receiving a drugs request. This message is labelled msg_3 in the diagram of Fig. 1 and is modeled using a shared variable $\texttt{msgVar}_3$ that is a parameter of PS. This event eventually triggers a transition in the process deliverySequence that corresponds to the end of the activity *drugsRequest* and also prompts the time observer process (via synchronization on port p_1) to start monitoring the elapsed time. The next state in line corresponds to the start of the *drugsChecking* activity that must be fulfilled within 6–12 unit of times unless the time observer enter its error state err (synchronization with the temporal observer on port p_2). At this point, the deliverySequence synchronizes on the "start" port of the <flow> process (port $\texttt{p}_3$ in Fig. 2). After the completion of the <flow> process (synchronization on port $\texttt{p}_4$) the deliverySequence gains hand again and the computation moves to the fulfillment of the reply activity *sendDrugsOrder*. Before concluding, the process interact again with the time observer (synchronization on port p_5) if the delay for sending the message $\texttt{msg}_4$ comply with the global constraint (c_3).

After presenting the BPEL transformation process to Fiacre language, let us present the real-time requirements we handle.

C. Defining real-Time Requirements

This section gives a description of the specification patterns used for the definition of real-time requirements available in our framework. A complete description of the language is given in [2]. Our language extends the property specification patterns of Dwyer et al. [14] with the ability to express time delays between the occurrences of events. The result is expressive enough to define a wide class of complex properties [2], [3]. The pattern language follows the classification introduced in [14], with patterns arranged in categories such as occurrence or order patterns. In the following, we study examples of *response* and *absence* patterns.

Absence pattern with delay.: This category of patterns can be used to specify delays within which activities must not occur. A typical pattern in this category can be used to assert that an activity, say A_2, cannot occur between d_1–d_2 units of time after the occurrence of an activity A_1. This requirement corresponds to a basic absence pattern in our language:

$$\texttt{absent}\ A_2\ \texttt{after}\ A_1\ \texttt{within}\,[d_1; d_2]\ . \qquad \text{(absent)}$$

An example of use for this pattern is the requirement that we cannot have two medical analyses for the patient in less than 10 days (240 hours):

$$\texttt{absent}\ \ MAS.medicalAnalysis\ \texttt{after}$$

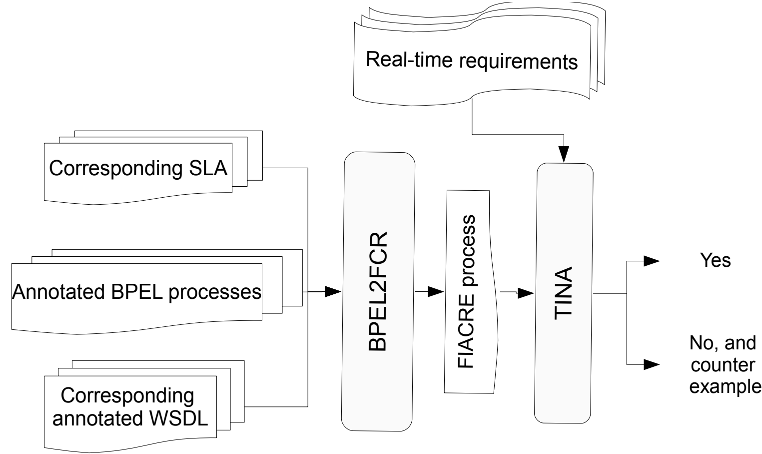

Figure 3. Bpel2Fcr: architecture of the transformation

$$MAS.medicalAnalysis \ \texttt{within} \ [0;\ 240]\ .$$

A more complicated example of requirement is to impose that if a doctor does not change a drug order within 6 hours, then it should not change drugs for another 48 hours. This requirement can be expressed using the composition of two absence patterns:

$$(\texttt{absent}\ MCS.drugsChanging\ \texttt{after}\ MCS.drugsAsking\ \texttt{within}\ [0;6]) \Rightarrow (\texttt{absent}\ MCS.drugsChanging\ \texttt{after}\ MCS.drugsAsking\ \texttt{within}\ [0;54])\ .$$

Response pattern with delay.: This category of patterns can be used to express delays between events, like for example constraints on the execution time of a service. The typical example of response pattern states that every occurrence of an event, say e_1, must be followed by an occurrence of an event e_2 within a time interval I. This pattern is denoted:

$$e_1\,\texttt{leadsto}\,e_2\,\texttt{within}\,I\ . \qquad \text{(leadsto)}$$

In our framework, we use the notation $S.init$ and $S.end$ to refer to a start, respectively an end, event in the service S. Hence, we can check that the execution time of the service S is less than d units of time with the requirement $S.init$ `leadsto` $S.end$ `within` $[0; d]$. We can use the same pattern to express requirements on an activity. For instance that drugs must be delivered within 48 hours of the medical examination start:

$$MCS.init\,\texttt{leadsto}\,PS.sendDrugsOrder\,\texttt{within}\,[0;48]\ .$$

Using a composition of patterns with the conjunction operator, we can bound the time between the start of the initiating service (the client) and the end of the choreography (if any). For instance, in our motivating example, we would like to check that the medical examination last less than 60 hours:

$$\begin{aligned} & MCS.init\,\texttt{leadsto}\,MCS.end\,\texttt{within}\,[0,60] \\ \wedge\ & MCS.init\,\texttt{leadsto}\,MAS.end\,\texttt{within}\,[0,60] \\ \wedge\ & MCS.init\,\texttt{leadsto}\,PS.end\,\texttt{within}\,[0,60]\ . \end{aligned}$$

More generally, we can check that a service, say S_2, always terminates its execution after service S_1 within a duration d with the requirement: $S_1.end$ `leadsto` $S_2.end$ `within` $[0; d]$

V. Experimentation

We have developed a prototype of a compiler from annotated BPEL processes to Fiacre in Java, named Bpel2Fcr, that is based on the interpretation defined in the previous section. The architecture of our transformation is depicted in Fig. 3. The input of our prototype is a set of BPEL processes, their corresponding WSDL and Service Level Agreement contracts that list the local and global constraints of the timed choreography. Our tool relies on EasyBPEL (http://easybpel.petalslink.org/), a library that provides a BPEL 2.0 engine to orchestrate services.

The Fiacre specification obtained as an output of Bpel2Fcr is the input of the model checker provided by the Tina verification toolbox to analyze specified choreography real-time requirements.

We give some results obtained with the analysis of our running example. The state graph for the HealthCare example has only 43 states and 51 transitions. This is obtained using optimized encoding (the graph for the unoptimized encoding is about three times as big). The generation of the Fiacre specification and its corresponding state space takes less than a second. For examples of this size, the verification time for checking a requirement is negligible; in the order of a couple of milliseconds.

This model is small due to the almost lack of concurrency in the scenario; only two `<flow>` activities and three

services. For more complex examples, we have tested our approach on an extended version of the healthcare scenario with seven different services. The resulting example has 886 states and 2476 transitions. With this larger example of timed choreography, the verification process takes in the order of half a second. As an example, we give the time and the validity of some real-time properties:

Property	*Result*	*Time (s)*
MCS.$init$ `leadsto` MCS.end `within` [0,60]	true	0,638
MCS.$init$ `leadsto` MCS.end `within` [0,20]	false	0,645
MAS.end `leadsto` MCS.end `within` [0,30]	true	0,531
PS.end `leadsto` MCS.end `within` [0,10]	false	0,637

VI. Conclusion

We describe a new framework for modeling and analyzing timed choreographies obtained through the composition of annotated BPEL processes. We most particularly focus on the problem of checking real-time requirements on a choreography. In this context, we have proposed a rich formal model for timed services composition that captures efficiently several kind of temporal constraints associated to the concurrent nature of services. The framework we propose has been implemented into a tool that automatically transforms timed BPEL processes into a Fiacre specification. In addition, we have shown an associated model-based verification approach to check real-time requirements on timed choreographies. In this context, we have defined classes of new real-time requirements which are specified using a logical-based formalism.

Work is still ongoing to improve and optimize the transformation implemented in our verification toolchain. Moreover, our ongoing work focuses on extending our approach by fault handling primitives and checking more complex requirements that relate to automatic reconfiguration of service composition. We should be assisted in that by the fact that our transformations are compositional.

References

[1] Cesar: Cost-efficient methods and processes for safety relevant embedded systems. *http://cesarproject.eu/*.

[2] N. Abid, S. Dal Zilio, and D. Le Botlan. Real-time specification patterns and tools. In *Proceedings of the International Workshop on Formal Methods for Industrial Critical Systems (FMICS'12), Lecture Notes in Computer Science, Vol. 7437,Springer-Verlag*, pages 1–15, 2012.

[3] N. Abid, S. Dal Zilio, and D. Le Botlan. A verified approach for checking real-time specification patterns. In *Proceedings of the 6th International Workshop on Verification and Evaluation of Computer and Communication Systems (VECoS'12)*, 2012.

[4] A. Airkin, S. Askary, S. Fordin, W. J. an D. Orchard, and K. Riemer. Web service choreography interface (wsci) 1.0. *W3C Working Group*, 2002.

[5] B. Benatallah, F. Casati, J. Ponge, and F. Toumani. On temporal abstractions of web service protocols. In *The 17th Conference on Advanced Information Systems Engineering (CAiSE '05). Short Paper Proceedings*, 2005.

[6] B. Benatallah, F. Casati, and F. Toumani. Analysis and management of web service protocols. *23rd International Conference on Conceptual Modeling*, November 2004.

[7] B. Benatallah, F. Casati, and F. Toumani. Web service conversation modeling: A cornerstone for e-business automation. *IEEE Internet Computing*, 8(1):46–54, 2004.

[8] B. Berthomieu, J.-P. Bodeveix, P. Farail, M. Filali, H. Garavel, P. Gauffilet, F. Lang, and F. Vernadat. Fiacre: an intermediate language for model verification in the topcased environment. In *Embedded Real Time Software (ERTS)*, 2008.

[9] B. Berthomieu, J.-P. Bodeveix, M. Filali, H. Garavel, F. Lang, F. Peres, R. Saad, J. Stoecker, and F. Vernadat. The syntax and semantics of fiacre. *Repport LAAS N 07264*, 2007.

[10] B. Berthomieu, P. O. Ribet, and F. Vernadat. The tool tina – construction of abstract state spaces for petri nets and time petri nets. *International Journal of Production Research*, 42(10), 2004.

[11] L. Bordeaux, G. Salaün, D. Berardi, and M. Mecella. When are two web services compatible? In *Technologies for E-Services, 5th International Workshop (TES)*, pages 15–28, 2004.

[12] M. D. Bozga, V. Sfyrla, and J. Sifakis. Modeling synchronous systems in bip. In *Proceedings of the seventh ACM international conference on Embedded software*, EMSOFT '09. ACM, 2009.

[13] T. Bultan and X. Fu. Choreography modeling and analysis with collaboration diagrams. *IEEE Data Eng. Bull.*, 31(3):27–30, 2008.

[14] M. B. Dwyer, G. S. Avrunin, and J. C. Corbett. Patterns in property specifications for finite-state verification. In *International Conference on Software Engineering (ICSE'99)*, pages 411–420, 1999.

[15] J. Eder and A. Tahamtan. Temporal conformance of federated choreographies. In *Proceedings of the* 19^{th} *International Conference on Database and Expert Systems Applications(DEXA'08)*, Turin, Italy, September 1-5, 2008.

[16] P. Farail, P. Gaufillet, A. Canals, C. Le Camus, D. Sciamma, P. Michel, X. Crégut, and M. Pantel. the TOPCASED project: a toolkit in open source for critical aeronautic systems design. In *Embedded Real Time Software (ERTS)*, 2006.

[17] N. Guermouche and C. Godart. Timed model checking based approach for web services analysis. In *IEEE International Conference on Web Services (ICWS'09), July 6-10, 2009, Los Angeles, CA, USA*, 2009.

[18] N. Guermouche and C. Godart. Asynchronous timed web service-aware choreography analysis. In *Proceedings of the* 21^{st} *International Conference on Information Systems Engineering (CAiSE'09)*, pages 364–378, Amsterdam, The Netherlands, June 8-12, 2009.

[19] N. Guermouche and C. Godart. Timed conversational protocol based approach for web services analysis. In *Proceedings of the* 8^{th} *International Conference on Service Oriented Computing (ICSOC'10)*, pages 603–611, San Francisco, California, USA, December 7-10, 2010.

[20] P.-C. Héam, O. Kouchnarenko, and J. Voinot. How to handle qos aspects in web services substitutivity verification. In *WETICE'07 – IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises*, pages 333–338, 2007.

[21] P.-C. Héam, O. Kouchnarenko, and J. Voinot. Component simulation-based substitutivity managing qos and composition issues. *Sci. Comput. Program.*, 75(10):898–917, 2010.

[22] S.-Y. Hwang, W.-F. Hsieh, and C.-H. Lee. Verifying web services in a choreography environment. In *Proceedings of the IEEE International Conference on Service-Oriented Computing and Applications (SOCA'11)*, pages 1–4, 2011.

[23] S. Kallel, A. Charfi, T. Dinkelaker, M. Mezini, and M. Jmaiel. Specifying and monitoring temporal properties in web services compositions. In *Proceedings of the* 7^{th} *IEEE European Conference on Web Services (ECOWS'09)*, pages 148–157, 2009.

[24] R. Kazhamiakin, P. K. Pandya, and M. Pistore. Representation, verification, and computation of timed properties in web service compositions. In *Proceedings of the IEEE International Conference on Web Services (ICWS)*, pages 497–504, 2006.

[25] R. Kazhamiakin, P. K. Pandya, and M. Pistore. Timed modelling and analysis in web service compositions. In *Proceedings of the The First International Conference on Availability, Reliability and Security, ARES*, pages 840–846. IEEE Computer Society, 2006.

[26] M. Makni, S. Tata, M. M. Yeddes, and N. B. Hadj-Alouane. Satisfaction and coherence of deadline constraints in inter-organizational workflows. In *Proceedings of the On the Move to Meaningful Internet Systems: OTM 2010 - Confederated International Conferences: CoopIS, IS, DOA and ODBASE,(OTM Conferences'10)*, pages 523–539, Hersonissos, Crete, Greece, October 25-29, 2010.

[27] P. M. Merlin and D. J. Farber. Recoverability of communication protocols: Implications of a theoretical study. *IEEE Trans. Comm.*, 24(9), 1976.

[28] J. Ponge, B. Benatallah, F. Casati, and F. Toumani. Fine-grained compatibility and replaceability analysis of timed web service protocols. In *the 26th International Conference on Conceptual Modeling (ER)*, 2007.